

DTIC COPY

Ada COMPILER
VALIDATION SUMMARY REPORT:
Certificate Number: 881121A1.10009
Alsys
AlsyCOMP_034, Version 4.1
Multitech 1100

Completion of On-Site Testing:
21 November 1988

Prepared By:
AFNOR
Tour Europe
Cedex 7
F-92080 Paris la Défense

Prepared For:
Ada Joint Program Office
United States Department of Defense
Washington DC 20301-3081

AD-A222 321

DTIC
ELECTE
MAY 17 1990
S B D

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

9-11-88 11 012

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	12. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Ada Compiler Validation Summary Report; Alsys, AlsyCOMP_034, Version 4.1, Multitech 1100 (Host & Target) 881121A1.10009		5. TYPE OF REPORT & PERIOD COVERED 21 Nov. 1988 to 21 Nov. 1989
7. AUTHOR(s) AFNOR, Paris, France.		6. PERFORMING ORG. REPORT NUMBER
8. PERFORMING ORGANIZATION AND ADDRESS AFNOR, Paris, France.		8. CONTRACT OR GRANT NUMBER(s)
11. CONTROLLING OFFICE NAME AND ADDRESS Ada Joint Program Office United States Department of Defense Washington, DC 20301-3081		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
14. MONITORING AGENCY NAME & ADDRESS (If different from Controlling Office) AFNOR, Paris, France.		12. REPORT DATE
		13. NUMBER OF PAGES
		15. SECURITY CLASS (of this report) UNCLASSIFIED
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE N/A
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20 if different from Report) UNCLASSIFIED		
18. SUPPLEMENTARY NOTES		
19. KEYWORDS (Continue on reverse side if necessary and identify by block number) Ada Programming language, Ada Compiler Validation Summary Report, Ada Compiler Validation Capability, ACVC, Validation Testing, Ada Validation Office, AVO, Ada Validation Facility, AVF, ANSI/MIL-STD- 1815A, Ada Joint Program Office, AJPO		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) Alsys, AlsyCOMP_034, Version 4.1, Paris La Defense, Multitech 1100 under Interactive 386/ix release 1.04 (Host & Target), ACVC 1.10.		

DD

FORM

1473

EDITION OF 1 NOV 85 IS OBSOLETE
S/N 0102-LF-014-0801

1 JAN 73

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

Ada Compiler Validation Summary Report:

Compiler Name: AlsyCOMP_034, Version 4.1

Certificate Number: 881121A1.10009

Host: Multitech 1100 under Interactive 386/ix release 1.04
(Unix system V.3)

Target: Multitech 1100 under Interactive 386/ix release 1.04
(Unix system V.3)

Testing Completed 21 November 1988 Using ACVC 1.10

This report has been reviewed and is approved.

F de Labareyre

AFNOR
Fabrice Garnier de Labareyre
Tour Europe
Cedex 7
F-92080 Paris la Défense

Dr. John F. Kramer
Ada Validation Organization
Dr. John F. Kramer
Institute for Defense Analyses
Alexandria VA 22311

Dr. John Solomond
Ada Joint Program Office
Dr. John Solomond
Director
Department of Defense
Washington DC 20301



Accession For	
NTIS GRA&I	<input checked="checked" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

TABLE OF CONTENTS

CHAPTER 1 INTRODUCTION

1.1 PURPOSE OF THIS VALIDATION SUMMARY REPORT5
1.2 USE OF THIS VALIDATION SUMMARY REPORT5
1.3 REFERENCES.6
1.4 DEFINITION OF TERMS6
1.5 ACVC TEST CLASSES7

CHAPTER 2 CONFIGURATION INFORMATION

2.1 CONFIGURATION TESTED.9
2.2 IMPLEMENTATION CHARACTERISTICS.	10

CHAPTER 3 TEST INFORMATION

3.1 TEST RESULTS.	15.
3.2 SUMMARY OF TEST RESULTS BY CLASS.	15
3.3 SUMMARY OF TEST RESULTS BY CHAPTER.	16
3.4 WITHDRAWN TESTS	16
3.5 INAPPLICABLE TESTS.	16
3.6 TEST, PROCESSING, AND EVALUATION MODIFICATIONS.	19
3.7 ADDITIONAL TESTING INFORMATION.	20
3.7.1 Prevalidation	20
3.7.2 Test Method	20
3.7.3 Test Site	21

APPENDIX A DECLARATION OF CONFORMANCE

APPENDIX B TEST PARAMETERS

APPENDIX C WITHDRAWN TESTS

APPENDIX D APPENDIX F OF THE Ada STANDARD



OFFICE OF THE DIRECTOR OF
DEFENSE RESEARCH AND ENGINEERING

WASHINGTON, DC 20301

19 APR 1990

(R&AT)

MEMORANDUM FOR Director, Directorate of Database Services,
Defense Logistics Agency

SUBJECT: Technology Screening of Unclassified/Unlimited Reports

Your letter of 2 February 1990 to the Commander, Air Force Systems Command, Air Force Aeronautical Laboratory, Wright-Patterson Air Force Base stated that the Ada Validation Summary report for Meridian Software Systems, Inc. contained technical data that should be denied public disclosure according to DoD Directive 5230.25

Feb 12
per
Asa
Easter

We do not agree with this opinion that the contents of this particular Ada Validation Summary Report or the contents of the several hundred of such reports produced each year to document the conformity testing results of Ada compilers. Ada is not used exclusively for military applications. The language is an ANSI Military Standard, a Federal Information Processing Standard, and an International Standards Organization standard. Compilers are tested for conformity to the standard as the basis for obtaining an Ada Joint Program Office certificate of conformity. The results of this testing are documented in a standard form in all Ada Validation Summary Reports which the compiler vendor agrees to make public as part of his contract with the testing facility.

On 18 December 1985, the Commerce Department issued Part 379 Technical Data of the Export Administration specifically listing Ada Programming Support Environments (including compilers) as items controlled by the Commerce Department. The AJPO complies with Department of Commerce export control regulations. When Defense Technical Information Center receives an Ada Validation Summary Report, which may be produced by any of the five U.S. and European Ada Validation Facilities, the content should be made available to the public.

If you have any further questions, please feel free to contact the undersigned at (202) 694-0209.

John P. Solomond

John P. Solomond
Director
Ada Joint Program Office

CHAPTER 1

INTRODUCTION

This Validation Summary Report (VSR) describes the extent to which a specific Ada compiler conforms to the Ada Standard, ANSI/MIL-STD-1815A. This report explains all technical terms used within it and thoroughly reports the results of testing this compiler using the Ada Compiler Validation Capability (ACVC). An Ada compiler must be implemented according to the Ada Standard, and any implementation-dependent features must conform to the requirements of the Ada Standard. The Ada Standard must be implemented in its entirety, and nothing can be implemented that is not in the Standard.

Even though all validated Ada compilers conform to the Ada Standard, it must be understood that some differences do exist between implementations. The Ada Standard permits some implementation dependencies--for example, the maximum length of identifiers or the maximum values of integer types. Other differences between compilers result from the characteristics of particular operating systems, hardware, or implementation strategies. All the dependencies observed during the process of testing this compiler are given in this report.

The information in this report is derived from the test results produced during validation testing. The validation process includes submitting a suite of standardized tests, the ACVC, as inputs to an Ada compiler and evaluating the results. The purpose of validating is to ensure conformity of the compiler to the Ada Standard by testing that the compiler properly implements legal language constructs and that it identifies and rejects illegal language constructs. The testing also identifies behavior that is implementation-dependent but permitted by the Ada Standard. Six classes of tests are used. These tests are designed to perform checks at compile time, at link time, and during execution.

1.1 PURPOSE OF THIS VALIDATION SUMMARY REPORT

This VSR documents the results of the validation testing performed on an Ada compiler. Testing was carried out for the following purposes:

- . To attempt to identify any language constructs supported by the compiler that do not conform to the Ada Standard
- . To attempt to identify any language constructs not supported by the compiler but required by the Ada Standard
- . To determine that the implementation-dependent behavior is allowed by the Ada Standard

Testing of this compiler was conducted under the direction of the AVF according to procedures established by the Ada Joint Program Office and administered by the Ada Validation Organization (AVO). On-site testing was completed 21 November 1988 at Alsys Inc. in Waltham, USA.

1.2 USE OF THIS VALIDATION SUMMARY REPORT

Consistent with the national laws of the originating country, the AVO may make full and free public disclosure of this report. In the United States, this is provided in accordance with the "Freedom of Information Act" (5 U.S.C.#552). The results of this validation apply only to the computers, operating systems, and compiler versions identified in this report.

The organizations represented on the signature page of this report do not represent or warrant that all statements set forth in this report are accurate and complete, or that the subject compiler has no nonconformities to the Ada Standard other than those presented. Copies of this report are available to the public from:

Ada Information Clearinghouse
Ada Joint Program Office
OUSDRE
The Pentagon, Rm 3D-139 (Fern Street)
Washington DC 20301-3081

or from:

AFNOR
Tour Europe
cedex 7
F-92080 Paris la Défense

Questions regarding this report or the validation test results should be directed to the AVF listed above or to:

Ada Validation Organization
Institute for Defense Analyses
1801 North Beauregard Street
Alexandria VA 22311

1.3 REFERENCES

1. Reference Manual for the Ada Programming Language, ANSI/MIL-STD-1815A, February 1983, and ISO 8652-1987.
2. Ada Compiler Validation Procedures and Guidelines, Ada Joint Program Office, 1 January 1987.
3. Ada Compiler Validation Capability Implementers' Guide, SofTech, Inc., December 1986.

1.4 DEFINITION OF TERMS

ACVC	The Ada Compiler Validation Capability. The set of Ada programs that tests the conformity of an Ada compiler to the Ada programming language.
Ada Commentary	An Ada Commentary contains all information relevant to the point addressed by a comment on the Ada Standard. These comments are given a unique identification number having the form AI-ddddd.
Ada Standard	ANSI/MIL-STD-1815A, February 1983 and ISO 8652-1987.
Applicant	The agency requesting validation.
AVF	The Ada Validation Facility. The AVF is responsible for conducting compiler validations according to procedures contained in the Ada Compiler Validation Procedures and Guidelines.
AVO	The Ada Validation Organization. The AVO has oversight authority over all AVF practices for the purpose of maintaining a uniform process for validation of Ada compilers. The AVO provides administrative and technical support for Ada validations to ensure consistent practices.
Compiler	A processor for the Ada language. In the context of this report, a compiler is any language processor,

including cross-compilers, translators, and interpreters.

Failed test	An ACVC test for which the compiler generates a result that demonstrates nonconformity to the Ada Standard.
Host	The computer on which the compiler resides.
Inapplicable test	An ACVC test that uses features of the language that a compiler is not required to support or may legitimately support in a way other than the one expected by the test.
Passed test	An ACVC test for which a compiler generates the expected result.
Target	The computer for which a compiler generates code.
Test	A program that checks a compiler's conformity regarding a particular feature or a combination of features to the Ada Standard. In the context of this report, the term is used to designate a single test, which may comprise one or more files.
Withdrawn test	An ACVC test found to be incorrect and not used to check conformity to the Ada Standard. A test may be incorrect because it has an invalid test objective, fails to meet its test objective, or contains illegal or erroneous use of the language.

1.5 ACVC TEST CLASSES

Conformity to the Ada Standard is measured using the ACVC. The ACVC contains both legal and illegal Ada programs structured into six test classes: A, B, C, D, E, and L. The first letter of a test name identifies the class to which it belongs. Class A, C, D, and E tests are executable, and special program units are used to report their results during execution. Class B tests are expected to produce compilation errors. Class L tests are expected to produce compilation or link errors.

Class A tests check that legal Ada programs can be successfully compiled and executed. There are no explicit program components in a Class A test to check semantics. For example, a Class A test checks that reserved words of another language (other than those already reserved in the Ada language) are not treated as reserved words by an Ada compiler. A Class A test is passed if no errors are detected at compile time and the program executes to produce a PASSED message.

Class B tests check that a compiler detects illegal language usage. Class B tests are not executable. Each test in this class is compiled and the resulting compilation listing is examined to verify that every syntax or semantic error in the test is detected. A Class B test is passed if every illegal construct that it contains is detected by the compiler.

Class C tests check that legal Ada programs can be correctly compiled and executed. Each Class C test is self-checking and produces a PASSED, FAILED, or NOT APPLICABLE message indicating the result when it is executed.

Class D tests check the compilation and execution capacities of a compiler. Since there are no capacity requirements placed on a compiler by the Ada Standard for some parameters--for example, the number of identifiers permitted in a compilation or the number of units in a library--a compiler may refuse to compile a Class D test and still be a conforming compiler. Therefore, if a Class D test fails to compile because the capacity of the compiler is exceeded, the test is classified as inapplicable. If a Class D test compiles successfully, it is self-checking and produces a PASSED or FAILED message during execution.

Each Class E test is self-checking and produces a NOT APPLICABLE, PASSED, or FAILED message when it is compiled and executed. However, the Ada Standard permits an implementation to reject programs containing some features addressed by Class E tests during compilation. Therefore, a Class E test is passed by a compiler if it is compiled successfully and executes to produce a PASSED message, or if it is rejected by the compiler for an allowable reason.

Class L tests check that incomplete or illegal Ada programs involving multiple, separately compiled units are detected and not allowed to execute. Class L tests are compiled separately and execution is attempted. A Class L test passes if it is rejected at link time--that is, an attempt to execute the main program must generate an error message before any declarations in the main program or any units referenced by the main program are elaborated.

Two library units, the package REPORT and the procedure CHECK_FILE, support the self-checking features of the executable tests. The package REPORT provides the mechanism by which executable tests report PASSED, FAILED, or NOT APPLICABLE results. It also provides a set of identity functions used to defeat some compiler optimizations allowed by the Ada Standard that would circumvent a test objective. The procedure CHECK_FILE is used to check the contents of text files written by some of the Class C tests for chapter 14 of the Ada Standard. The operation of REPORT and CHECK_FILE is checked by a set of executable tests. These tests produce messages that are examined

to verify that the units are operating correctly. If these units are not operating correctly, then the validation is not attempted.

The text of the tests in the ACVC follow conventions that are intended to ensure that the tests are reasonably portable without modification. For example, the tests make use of only the basic set of 55 characters, contain lines with a maximum length of 72 characters, use small numeric values, and place features that may not be supported by all implementations in separate tests. However, some tests contain values that require the test to be customized according to implementation-specific values--for example, an illegal file name. A list of the values used for this validation is provided in Appendix C.

A compiler must correctly process each of the tests in the suite and demonstrate conformity to the Ada Standard by either meeting the pass criteria given for the test or by showing that the test is inapplicable to the implementation. The applicability of a test to an implementation is considered each time the implementation is validated. A test that is inapplicable for one validation is not necessarily inapplicable for a subsequent validation. Any test that was determined to contain an illegal language construct or an erroneous language construct is withdrawn from the ACVC and, therefore, is not used in testing a compiler. The tests withdrawn at the time of this validation are given in Appendix D.

CHAPTER 2

CONFIGURATION INFORMATION

2.1 CONFIGURATION TESTED

The candidate compilation system for this validation was tested under the following configuration:

Compiler: AlsyCOMP_034, Version 4.1

ACVC Version: 1.10

Certificate Number: 881121A1.10009

Host Computer:

Machine: Multitech 1100

Operating System: Unix
Interactive 386/ix release 1.04
(Unix V.3)

Memory Size: 9 Mb

Configuration information :

monochrome display and adapter
machine 1: 40 Mb hard disk
machine 2: 100 Mb hard disk
80387 mathematic
co-processor

Target Computer:

Machine: Multitech 1100

Operating System: Unix
Interactive 386/ix release 1.04
(Unix V.3)

Memory Size: 9 Mb

Configuration information :

monochrome display and adapter
 machine 1: 40 Mb hard disk
 machine 2: 100 Mb hard disk
 80387 mathematic
 co-processor

Communications Network: Micom board for connexion to a Vax
 for file transfer

2.2 IMPLEMENTATION CHARACTERISTICS

One of the purposes of validating compilers is to determine the behavior of a compiler in those areas of the Ada Standard that permit implementations to differ. Class D and E tests specifically check for such implementation differences. However, tests in other classes also characterize an implementation. The tests demonstrate the following characteristics:

. Capacities.

The compiler correctly processes a compilation containing 723 variables in the same declarative part. (See test D29002K.)

The compiler correctly processe a test containing loop statements nested to 65 levels. (See tests D55A03A..H (8 tests).)

The compiler correctly processes tests containing block statements nested to 65 levels. (See test D56001B.)

The compiler correctly processes tests containing recursive procedures separately compiled as subunits nested to 17 levels. (See tests D64005E..G (3 tests).)

. Predefined types.

This implementation supports the additional predefined types SHORT_INTEGER, SHORT_SHORT_INTEGER, LONG_FLOAT in the package STANDARD. (See tests B86001T..Z (7 tests).)

. Based literals.

An implementation is allowed raise `NUMERIC_ERROR` or `CONSTRAINT_ERROR` when a value exceeds `SYSTEM.MAX_INT`. This implementation raises `CONSTRAINT_ERROR` during execution. (See test E24201A.)

. Expression evaluation.

Apparently no default initialization expressions for record components are evaluated before any value is checked to belong to a component's subtype. (See test C32117A.)

Assignments for subtypes are performed with the same precision as the base type. (See test C35712B.)

This implementation uses no extra bits for extra precision. This implementation uses all extra bits for extra range. (See test C35903A.)

Apparently `NUMERIC_ERROR` is raised when an integer literal operand in a comparison or membership test is outside the range of the base type. (See test C45232A.)

Apparently `NUMERIC_ERROR` is raised when a literal operand in a fixed-point comparison or membership test is outside the range of the base type. (See test C45252A.)

Apparently underflow is gradual. (See tests C45524A..Z.)

. Rounding.

The method used for rounding to integer is apparently round to even. (See tests C46012A..Z.)

The method used for rounding to longest integer is apparently round to even. (See tests C46012A..Z.)

The method used for rounding to integer in static universal real expressions is apparently round to even. (See test C4A014A.)

. Array types.

An implementation is allowed to raise `NUMERIC_ERROR` or `CONSTRAINT_ERROR` for an array having a `'LENGTH` that exceeds

STANDARD.INTEGER'LAST and/or SYSTEM.MAX_INT. For this implementation:

Declaration of an array type or subtype declaration with more than SYSTEM.MAX_INT components raises NUMERIC_ERROR sometimes, CONSTRAINT_ERROR sometimes. (See test C36003A.)

CONSTRAINT_ERROR is raised when 'LENGTH is applied to an array type with INTEGER'LAST + 2 components. (See test C36202A.)

CONSTRAINT_ERROR is raised when an array type with SYSTEM.MAX_INT + 2 components is declared. (See test C36202B.)

A packed BOOLEAN array having a 'LENGTH exceeding INTEGER'LAST raises no exception. (See test C52103X.)

A packed two-dimensional BOOLEAN array with more than INTEGER'LAST components raises CONSTRAINT_ERROR when the length of a dimension is calculated and exceeds INTEGER'LAST. array objects are sliced. (See test C52104Y.)

A null array with one dimension of length greater than INTEGER'LAST may raise NUMERIC_ERROR or CONSTRAINT_ERROR either when declared or assigned. Alternatively, an implementation may accept the declaration. However, lengths must match in array slice assignments. This implementation raises no exception. (See test E52103Y.)

In assigning one-dimensional array types, the expression appears to be evaluated in its entirety before CONSTRAINT_ERROR is raised when checking whether the expression's subtype is compatible with the target's subtype. (See test C52013A.)

In assigning two-dimensional array types, the expression does not appear to be evaluated in its entirety before CONSTRAINT_ERROR is raised when checking whether the expression's subtype is compatible with the target's subtype. (See test C52013A.)

Discriminated types.

In assigning record types with discriminants, the expression appears to be evaluated in its entirety before CONSTRAINT_ERROR is raised when checking whether the expression's subtype is compatible with the target's subtype. (See test C52013A.)

. Aggregates.

In the evaluation of a multi-dimensional aggregate, all choices appear to be evaluated before checking against the index type. (See tests C43207A and C43207B.)

In the evaluation of an aggregate containing subaggregates, not all choices are evaluated before being checked for identical bounds. (See test E43212B.)

All choices are evaluated before CONSTRAINT_ERROR is raised if a bound in a non-null range of a non-null aggregate does not belong to an index subtype. (See test E43211B.)

. Pragma.

The pragma INLINE is supported for functions or procedures, but not functions called inside a package specification. (See tests LA3004A..B, EA3004C..D, and CA3004E..F.)

Pragma PACK is not supported for records. (See test ED1D04A.)

. Generics.

Generic specifications and bodies can be compiled in separate compilations. (See tests CA1012A, CA2009C, CA2009F, BC3204C, and BC3205D.)

Generic unit bodies and their subunits can be compiled in separate compilations. (See test CA3011A.)

Generic subprogram declarations and bodies can be compiled in separate compilations. (See tests CA1012A and CA2009F.)

Generic library subprogram specifications and bodies can be compiled in separate compilations. (See test CA1012A.)

Generic non-library subprogram bodies can be compiled in separate compilations from their stubs. (See test CA2009F.)

Generic package declarations and bodies can be compiled in separate compilations. (See tests CA2009C, BC3204C, and BC3205D.)

Generic library package specifications and bodies can be compiled in separate compilations. (See tests BC3204C and BC3205D.)

Generic non-library package bodies as subunits can be compiled in separate compilations. (See test CA2009C.)

Generic unit bodies and their subunits can be compiled in separate compilations. (See test CA3011A.)

Input and output.

The package SEQUENTIAL_IO can be instantiated with unconstrained array types and record types with discriminants without defaults. (See tests AE2101C, EE2201D, and EE2201E.)

The package DIRECT_IO can not be instantiated with unconstrained array types and record types with discriminants without defaults. (See tests AE2101H, EE2401D, and EE2401G.)

Modes IN_FILE and OUT_FILE are supported for SEQUENTIAL_IO, but not CREATE in mode IN_FILE. (See tests CE2102D..E, CE2102N, and CE2102P.)

Modes IN_FILE, OUT_FILE, and INOUT_FILE are supported for DIRECT_IO, but not CREATE in mode IN_FILE. (See tests CE2102F, CE2102I..J, CE2102R, CE2102T, and CE2102V.)

Modes IN_FILE and OUT_FILE are supported for text files, but not CREATE in mode IN_FILE. (See tests CE3102E and CE3102I..K.)

RESET from OUT_FILE to IN_FILE only and DELETE operations are supported for SEQUENTIAL_IO. (See tests CE2102G and CE2102X.)

RESET except from IN_FILE to INOUT_FILE or to OUT_FILE and DELETE operations are supported for DIRECT_IO. (See tests CE2102K and CE2102Y.)

RESET and DELETE operations are supported for text files. (See tests CE3102F..G, CE3104C, CE3110A, and CE3114A.)

Overwriting to a sequential file truncates to the last element written. (See test CE2208B.)

Temporary sequential files are given names and deleted when closed. (See test CE2108A.)

Temporary direct files are given names and deleted when closed. (See test CE2108C.)

Temporary text files are given names and deleted when closed. (See test CE3112A.)

More than one internal file can be associated with each external file for sequential files when reading only (See tests CE2107A..E, CE2102L, CE2110B, and CE2111D.)

More than one internal file can be associated with each external file for direct files when reading only (See tests CE2107F..I, CE2110D and CE2111H.)

More than one internal file can be associated with each external file for text files when reading only. (See tests CE3111A..E, CE3114B, and CE3115A.)

CHAPTER 3

TEST INFORMATION

3.1 TEST RESULTS

Version 1.10 of the ACVC comprises 3717 tests. When this compiler was tested, 36 tests had been withdrawn because of test errors. The AVF determined that 387 tests were inapplicable to this implementation. All inapplicable tests were processed during validation testing except for 201 executable tests that use floating-point precision exceeding that supported by the implementation. Modifications to the code, processing, or grading for 52 tests were required to successfully demonstrate the test objective. (See section 3.6.)

The AVF concludes that the testing results demonstrate acceptable conformity to the Ada Standard.

3.2 SUMMARY OF TEST RESULTS BY CLASS

RESULT	TEST CLASS						TOTAL
	A	B	C	D	E	L	
Passed	129	1130	1940	17	32	46	3294
Inapplicable	0	8	377	0	2	0	387
Withdrawn	1	2	33	0	0	0	36
TOTAL	130	1140	2350	17	34	46	3717

3.3 SUMMARY OF TEST RESULTS BY CHAPTER

RESULT	CHAPTER														TOTAL
	2	3	4	5	6	7	8	9	10	11	12	13	14		
Passed	199	577	545	245	172	99	161	332	137	36	252	259	280	3292	
Inappl	14	72	135	3	0	0	5	1	0	0	0	116	41	387	
Wdrn	0	1	0	0	0	0	0	1	0	0	1	29	4	36	
TOTAL	213	650	680	248	172	99	166	334	137	36	253	404	325	3717	

3.4 WITHDRAWN TESTS

The following 36 tests were withdrawn from ACVC Version 1.10 at the time of this validation:

A39005G	B97102E	BC3009B	CD2A62D	CD2A63A	CD2A63B	CD2A63C
CD2A63D	CD2A66A	CD2A66B	CD2A66C	CD2A66D	CD2A73A	CD2A73B
CD2A73C	CD2A73D	CD2A76A	CD2A76B	CD2A76C	CD2A76D	CD2A81G
CD2A83G	CD2A84N	CD2A84M	CD50110	CD2B15C	CD7205C	CD5007B
CD7105A	CD7203B	CD7204B	CD7205D	CE2107I	CE3111C	CE3301A
CE3411B						

See Appendix D for the reason that each of these tests was withdrawn.

3.5 INAPPLICABLE TESTS

Some tests do not apply to all compilers because they make use of features that a compiler is not required by the Ada Standard to support. Others may depend on the result of another test that is either inapplicable or withdrawn. The applicability of a test to an implementation is considered each time a validation is attempted. A test that is inapplicable for one validation attempt is not necessarily inapplicable for a subsequent attempt. For this validation attempt, 387 tests were inapplicable for the reasons indicated:

The following 201 tests are not applicable because they have floating-point type declarations requiring more digits than System.Max_Digits:

C24113L..Y	C35705L..Y	C35706L..Y	C35707L..Y	C35708L..Y
C35802L..Z	C45241L..Y	C45321L..Y	C45421L..Y	C45521L..Z
C45524L..Z	C45621L..Z	C45641L..Y	C46012L..Z	

- . C35702A and B86001T are not applicable because this implementation supports no predefined type Short_Float.
- . C45531M..P (4 tests) and C45532M..P (4 tests) are not applicable because the value of System.Max_Mantissa is less than 32.
- . The following 16 tests are not applicable because this implementation does not support a predefined type LONG_INTEGER:

C45231C	C45304C	C45502C	C45503C	C45504C
C45504F	C45611C	C45613C	C45614C	C45631C
C45632C	B52004D	C55B07A	B55B09C	B86001W
CD7101F				
- . C86001F, is not applicable because recompilation of Package SYSTEM is not allowed.
- . B86001Y is not applicable because this implementation supports no predefined fixed-point type other than Duration.
- . B86001Z is not applicable because this implementation supports no predefined floating-point type with a name other than Float, Long_Float, or Short_Float.
- . B91001H is not applicable because address clause for entries is not supported by this implementation.
- . CD1009C, CD2A41A..B, CD2A41E, CD2A42A..B, CD2A42E..F, CD2A42I..J are not applicable because size clause on float is not supported by this implementation.
- . CD1C04B, CD1C04E, CD4051A..D are not applicable because representation clause on derived records or derived tasks is not supported by this implementation.
- . CD1C03C, CD2A83A..C, CD2A83E, CD2A84B..I, CD2A84K..L, CD2B11B are not applicable because storage size clause on collection of unconstrained object is not supported.
- . CD1C04A, CD2A21C..D, CD2A22C..D, CD2A22G..H, CD2A31C..D, CD2A32C..D, CD2A32G..H, CD2A41C..D, CD2A42C..D, CD2A42G..H, CD2A51C..D, CD2A52C..D, CD2A52G..H, CD2A53D, CD2A54D, CD2A54H are not applicable because size clause for derived private type is not supported by this implementation.
- . CD2A61A..D,F,H,I,J,K,L, CD2A62A..C, CD2A71A..D, CD2A72A..D, CD2A74A..D, CD2A75A..D are not applicable because of the way this implementation allocates storage space for one component, size specification clause for an array type or for a record type

requires compression of the storage space needed for all the components (without gaps).

- . CD4041A is not applicable because alignment "at mod 8" is not supported by this implementation.
- . BD5006D is not applicable because address clause for packages is not supported by this implementation.
- . CD5011B,D,F,H,L,N,R, CD5012C,D,G,H,L, CD5013B,D,F,H,L,N,R, CD5014U,W are not applicable because address clause for a constant is not supported by this implementation.
- . CD5012J, CD5013S, CD5014S are not applicable because address clause for a task is not supported by this implementation.
- . CE2102E is inapplicable because this implementation supports create with out_file mode for SEQUENTIAL_IO.
- . CE2102F is inapplicable because this implementation supports create with inout_file mode for DIRECT_IO.
- . CE2102J is inapplicable because this implementation supports create with out_file mode for DIRECT_IO.
- . CE2102N is inapplicable because this implementation supports open with in_file mode for SEQUENTIAL_IO.
- . CE2102O is inapplicable because this implementation supports RESET with in_file mode for SEQUENTIAL_IO.
- . CE2102P is inapplicable because this implementation supports open with out_file mode for SEQUENTIAL_IO.
- . CE2102Q is inapplicable because this implementation supports RESET with out_file mode for SEQUENTIAL_IO.
- . CE2102R is inapplicable because this implementation supports open with inout_file mode for DIRECT_IO.
- . CE2102S is inapplicable because this implementation supports RESET with inout_file mode for DIRECT_IO.
- . CE2102T is inapplicable because this implementation supports open with in_file mode for DIRECT_IO.
- . CE2102U is inapplicable because this implementation supports RESET with in_file mode for DIRECT_IO.

- . CE2102V is inapplicable because this implementation supports open with out_file mode for DIRECT_IO.
- . CE2102W is inapplicable because this implementation supports RESET with out_file mode for DIRECT_IO.
- . CE2105A is not applicable because create with mode in_file is not supported by this implementation for SEQUENTIAL_IO.
- . CE2105B is inapplicable because CREATE with IN_FILE mode is not supported for direct access files.
- . CE2107B..E (4 tests), CE2107L, and CE2110B are not applicable because multiple internal files cannot be associated with the same external file when one or more files is writing for sequential files. The proper exception is raised when multiple access is attempted.
- . CE2107G..H (2 tests), CE2110D, and CE2111H are not applicable because multiple internal files cannot be associated with the same external file when one or more files is writing for direct files. The proper exception is raised when multiple access is attempted.
- . CE2111C,D are not applicable because resetting from in_file to out_file mode for sequential files is not supported by this implementation.
- . EE2401D and EE2401G are not applicable because USE_ERROR is raised when the create of an instantiation of DIRECT_IO with unconstrained type is called.
- . CE2401H is not applicable because create with inout_file mode for unconstrained records with default discriminants is not supported by this implementation.
- . CE3102F is inapplicable because this implementation supports reset for text files, for out_file, in_file and from out_file to in_file mode.
- . CE3102G is inapplicable because this implementation supports deletion of an external file for text files.
- . CE3102I is inapplicable because this implementation supports create with out_file mode for text files.
- . CE3102J is inapplicable because this implementation supports open with in_file mode for text files.

- . CE3102K is inapplicable because this implementation supports open with out_file mode for text files.
- . CE3109A is inapplicable because text file create with in_file mode is not supported and raises USE_ERROR.
- . CE3111B, CE3111D..E (2 tests), CE3114B, and CE3115A are not applicable because multiple internal files cannot be associated with the same external file when one or more files is writing for text files. The proper exception is raised when multiple access is attempted.

3.6 TEST, PROCESSING, AND EVALUATION MODIFICATIONS

It is expected that some tests will require modifications of code, processing, or evaluation in order to compensate for legitimate implementation behavior. Modifications are made by the AVF in cases where legitimate implementation behavior prevents the successful completion of an (otherwise) applicable test. Examples of such modifications include: adding a length clause to alter the default size of a collection; splitting a Class B test into subtests so that all errors are detected; and confirming that messages produced by an executable test demonstrate conforming behavior that wasn't anticipated by the test (such as raising one exception instead of another).

Modifications were required for 52 tests.

The test EA3004D when run as it is, the implementation fails to detect an error on line 27 of test file EA3004D6M (line 115 of "cat -n ea3004d*"). This is because the pragma INLINE has no effect when its object is within a package specification. However, the results of running the test as it is does not confirm that the pragma had no effect, only that the package was not made obsolete. By re-ordering the compilations so that the two subprograms are compiled after file D5 (the re-compilation of the "with"ed package that makes the various earlier units obsolete), we create a test that shows that indeed pragma INLINE has no effect when applied to a subprogram that is called within a package specification: the test then executes and produces the expected NOT_APPLICABLE result (as though INLINE were not supported at all). The re-ordering of EA3004D test files is 0-1-4-5-2-3-6.

The following 30 tests were split because syntax errors at one point resulted in the compiler not detecting other errors in the test:

B23004A	B24007A	B24009A	B25002A	B26005A	B27005A	B28003A
B32202A	B32202B	B32202C	B33001A	B36307A	B37004A	B49003A
B49005A	B61012A	B62001B	B74304B	B74304C	B74401F	B74401R
B91004A	B95032A	B95069A	B95069B	BA1101B2	BA1101B4	BC2001D

BC3009A BC3009C BD5005B

The following 21 tests were split in order to show that the compiler was able to find the representation clause indicated by the comment --N/A =>ERROR :

CD2A61A	CD2A61B	CD2A61F	CD2A61I	CD2A61J	CD2A62A	CD2A62B
CD2A71A	CD2A71B	CD2A72A	CD2A72B	CD2A75A	CD2A75B	CD2A84B
CD2A84C	CD2A84D	CD2A84E	CD2A84F	CD2A84G	CD2A84H	CD2A84I

3.7 ADDITIONAL TESTING INFORMATION

3.7.1 Prevalidation

Prior to validation, a set of test results for ACVC Version 1.10 produced by the AlsyCOMP_034 was submitted to the AVF by the applicant for review. Analysis of these results demonstrated that the compiler successfully passed all applicable tests, and the compiler exhibited the expected behavior on all inapplicable tests.

3.7.2 Test Method

Testing of the AlsyCOMP_034 using ACVC Version 1.10 was conducted on-site by a validation team from the AVF. The configuration consisted of a Multitech 1100 operating under Interactive 386/ix release 1.04 (Unix system V.3).

A tape containing all tests was taken on-site by the validation team for processing. Tests that make use of implementation-specific values were customized by Alsys after loading of the tape.

The contents of the tape were not loaded directly onto the host computer. They were loaded on a VAX/VMS machine and transferred via a network to the Multitech 1100. This is the reason why prevalidation tests were used for the the validation. Those tests were loaded by Alsys from a magnetic tape containing all tests provided by the AVF. Customization was done by Alsys. All the tests were checked at prevalidation time.

Integrity of the validation tests was made by checking that no modification of the test occurred after the time the prevalidation results were transferred on disquettes for submission to the AVF. This check was performed by verifying that the date of creation (or last modification) of the test files was earlier than the

prevalidation date. After validation was performed, 80 source tests were selected by the AVF and checked for integrity.

The full set of tests was compiled, linked, and all executable tests were run on the Multitech 1100. Results were printed from the from the host computer.

The compiler was tested using command scripts provided by Alsys and reviewed by the validation team. The compiler was tested using all default option settings except for the following:

OPTION / SWITCH EFFECT

CALLS=INLINE The pragma INLINE are taken into account

Tests were compiled, linked, and executed (as appropriate) using a single computer. Test output, compilation listings, and job logs were captured on disquettes and archived at the AVF. The listings examined on-site by the validation team were also archived.

3.7.3 Test Site

Testing was conducted at Alsys, Inc. in Waltham, USA and was completed on 21 November 1988.

Storage error occurred during processing of executable tests. Compilation of tests were restarted from the beginning of the chapter.

APPENDIX A

DECLARATION OF CONFORMANCE

Alsys has submitted the following Declaration of
Conformance concerning the AlsyCOMP_034.

DECLARATION OF CONFORMANCE

Compiler Implementor: Alsys

Ada Validation Facility: AFNOR, Tour Europe Cedex 7,
F-92080 Paris la Défense

Ada Compiler Validation Capability (ACVC) Version: 1.10

Base Configuration

Base Compiler Name: AlsyCOMP_034 Version 4.1

Host Architecture ISA: Multitech 1100
OS&VER #: Interactive 386/ix release 1.04 (Unix
system V.3)

Target Architecture ISA: Multitech 1100
OS&VER #: Interactive 386/ix release 1.04 (Unix
system V.3)

Implementor's Declaration

I, the undersigned, representing Alsys, have implemented no deliberate extensions to the Ada Language Standard ANSI/MIL-STD-1815A in the compiler(s) listed in this declaration. I declare that Alsys is the owner of record of the Ada language compiler(s) listed above and, as such, is responsible for maintaining said compiler(s) in conformance to ANSI/MIL-STD-1815A. All certificates and registrations for Ada language compiler(s) listed in this declaration shall be made only in the owner's corporate name.

Alsys
Mike Blanchette
Vice President and Director of Engineering

Date _____

Owner's Declaration

I, the undersigned, representing Alsys, take full responsibility for implementation and maintenance of the Ada compiler(s) listed above, and agree to the public disclosure of the final Validation Summary Report. I further agree to continue to comply with the Ada trademark policy, as defined by the Ada Joint Program Office. I declare that all of the Ada language compilers listed, and their host/target performance, are in compliance with the Ada Language Standard ANSI/MIL-STD-1815A.

Alsys
Mike Blanchette
Vice President and Director of Engineering

Date _____

APPENDIX B

TEST PARAMETERS

Certain tests in the ACVC make use of implementation-dependent values, such as the maximum length of an input line and invalid file names. A test that makes use of such values is identified by the extension .TST in its file name. Actual values to be substituted are represented by names that begin with a dollar sign. A value must be substituted for each of these names before the test is run. The values used for this validation are given below.

Name and Meaning	Value
-----	-----
\$ACC_SIZE An integer literal whose value is the number of bits sufficient to hold any value of an access type.	32
\$BIG_ID1 Identifier the size of the maximum input line length with varying last character.	(254 * 'A') & '1'
\$BIG_ID2 Identifier the size of the maximum input line length with varying last character.	(254 * 'A') & '2'
\$BIG_ID3 Identifier the size of the maximum input line length with varying middle character.	(126 * 'A') & '3' & (128 * 'A')
\$BIG_ID4 Identifier the size of the maximum input line length with varying middle character.	(126 * 'A') & '4' & (128 * 'A')

Name and Meaning	Value
\$BIG_INT_LIT An integer literal of value 298 with enough leading zeroes so that it is the size of the maximum line length.	(252 * '0') & '298'
\$BIG_REAL_LIT A universal real literal of value 690.0 with enough leading zeroes to be the size of the maximum line length.	(250 * '0') & '690.0'
\$BIG_STRING1 A string literal which when catenated with BIG_STRING2 yields the image of BIG_ID1.	' ' & (127 * 'A') & ' '
\$BIG_STRING2 A string literal which when catenated to the end of BIG_STRING1 yields the image of BIG_ID1.	' ' & (127 * 'A') & "1"
\$BLANKS A sequence of blanks twenty characters less than the size of the maximum line length.	(235 * ' ')
\$COUNT_LAST A universal integer literal whose value is TEXT_IO.COUNT'LAST.	214783647.
\$DEFAULT_MEM_SIZE An integer literal whose value is SYSTEM.MEMORY_SIZE.	655360
\$DEFAULT_STOR_UNIT An integer literal whose value is SYSTEM.STORAGE_UNIT.	8
\$DEFAULT_SYS_NAME The value of the constant SYSTEM.SYSTEM_NAME.	I_80X86
\$DELTA_DOC A real literal whose value is SYSTEM.FINE_DELTA.	2#1.0#E-31

Name and Meaning	Value
\$FIELD_LAST A universal integer literal whose value is TEXT_IO.FIELD'LAST.	255
\$FIXED_NAME The name of a predefined fixed-point type other than DURATION.	NO_SUCH_FIXED_TYPE
\$FLOAT_NAME The name of a predefined floating-point type other than FLOAT, SHORT_FLOAT, or LONG_FLOAT.	NO_SUCH_TYPE
\$GREATER_THAN_DURATION A universal real literal that lies between DURATION'BASE'LAST and DURATION'LAST or any value in the range of DURATION.	2_097_151.999_023_437_51
\$GREATER_THAN_DURATION_BASE_LAST A universal real literal that is greater than DURATION'BASE'LAST.	3_000_000.0
\$HIGH_PRIORITY An integer literal whose value is the upper bound of the range for the subtype SYSTEM.PRIORITY.	10
\$ILLEGAL_EXTERNAL_FILE_NAME1 An external file name which contains invalid characters.	ILLEGAL\!\$%`&*()/_+~
\$ILLEGAL_EXTERNAL_FILE_NAME2 An external file name which is too long.	!\$%`&*()?)/*&\!\$%`
\$INTEGER_FIRST A universal integer literal whose value is INTEGER'FIRST.	-32768
\$INTEGER_LAST A universal integer literal whose value is INTEGER'LAST.	32767

Name and Meaning	Value
-----	-----
\$INTEGER_LAST_PLUS_1 A universal integer literal whose value is INTEGER'LAST + 1.	32768
\$LESS_THAN_DURATION A universal real literal that lies between DURATION'BASE'FIRST and DURATION'FIRST or any value in the range of DURATION.	-2_097_152.5
\$LESS_THAN_DURATION_BASE_FIRST A universal real literal that is less than DURATION'BASE'FIRST.	-3000_000.0
\$LOW_PRIORITY An integer literal whose value is the lower bound of the range for the subtype SYSTEM.PRIORITY.	1
\$MANTISSA_DOC An integer literal whose value is SYSTEM.MAX_MANTISSA.	31
\$MAX_DIGITS Maximum digits supported for floating-point types.	15
\$MAX_IN_LEN Maximum input line length permitted by the implementation.	255
\$MAX_INT A universal integer literal whose value is SYSTEM.MAX_INT.	2147483647
\$MAX_INT_PLUS_1 A universal integer literal whose value is SYSTEM.MAX_INT+1.	2147483648
\$MAX_LEN_INT_BASED_LITERAL A universal integer based literal whose value is 2#11# with enough leading zeroes in the mantissa to be MAX_IN_LEN long.	'42:' & (250 * '0') & '11:'

Name and Meaning	Value
\$MAX_LEN_REAL_BASED_LITERAL A universal real based literal whose value is 16: F.E: with enough leading zeroes in the mantissa to be MAX_IN_LEN long.	'16:' & (248 * '0') & 'F.E:'
\$MAX_STRING_LITERAL A string literal of size MAX_IN_LEN, including the quote characters.	''' & (253 * 'A') & '''
\$MIN_INT A universal integer literal whose value is SYSTEM.MIN_INT.	-2147483648
\$MIN_TASK_SIZE An integer literal whose value is the number of bits required to hold a task object which has no entries, no declarations, and NULL;" as the only statement in its body.	32
\$NAME A name of a predefined numeric type other than FLOAT, INTEGER, SHORT_FLOAT, SHORT_INTEGER, LONG_FLOAT, or LONG_INTEGER.	NO_SUCH_TYPE
\$NAME_LIST A list of enumeration literals in the type SYSTEM.NAME, separated by commas.	I_80X86
\$NEG_BASED_INT A based integer literal whose highest order nonzero bit falls in the sign bit position of the representation for SYSTEM.MAX_INT.	16#FFFFFFFFE#
\$NEW_MEM_SIZE An integer literal whose value is a permitted argument for pragma memory_size, other than DEFAULT_MEM_SIZE. If there is no other value, then use DEFAULT_MEM_SIZE.	655360

Name and Meaning	Value
-----	-----
\$NEW_STOR_UNIT An integer literal whose value is a permitted argument for pragma storage_unit, other than DEFAULT_STOR_UNIT. If there is no other permitted value, then use value of SYSTEM.STORAGE_UNIT.	8
\$NEW_SYS_NAME A value of the type SYSTEM.NAME, other than \$DEFAULT_SYS_NAME. If there is only one value of that type, then use that value.	I_80X86
\$TASK_SIZE An integer literal whose value is the number of bits required to hold a task object which has a single entry with one inout parameter.	32
\$TICK A real literal whose value is SYSTEM.TICK.	1.0/18.2

APPENDIX C

WITHDRAWN TESTS

Some tests are withdrawn from the ACVC because they do not conform to the Ada Standard. The following 36 tests had been withdrawn at the time of validation testing for the reasons indicated. A reference of the form AI-ddddd is to an Ada Commentary.

- A39005G This test unreasonably expects a component clause to pack an array component into a minimum size (line 30).
- B97102E This test contains an unintended illegality: a select statement contains a null statement at the place of a selective wait alternative (line 31).
- BC3009B This test wrongly expects that circular instantiations will be detected in several compilation units even though none of the units is illegal with respect to the units it depends on; by AI-00256, the illegality need not be detected until execution is attempted (line 95).
- CD2A62D This test wrongly requires that an array object's size be no greater than 10 although its subtype's size was specified to be 40 (line 137).
- CD2A63A..D, CD2A66A..D, CD2A73A..D, CD2A76A..D [16 tests] These tests wrongly attempt to check the size of objects of a derived type (for which a 'SIZE length clause is given) by passing them to a derived subprogram (which implicitly converts them to the parent type (Ada standard 3.4:14)). Additionally, they use the 'SIZE length clause and attribute, whose interpretation is considered problematic by the WG9 ARG.
- CD2A81G, CD2A83G, CD2A84N & M, & CD50110 These tests assume that dependent tasks will terminate while the main program executes a loop that simply tests for task termination; this is not the case, and the main program may loop indefinitely (lines 74, 85, 86 & 96, 86 & 96, and 58, resp.).
- CD2B15C & CD7205C These tests expect that a 'STORAGE_SIZE length clause provides precise control over the number of designated objects in a collection; the Ada standard 13.2:15 allows that such control must not be expected.

- CD5007B This test wrongly expects an implicitly declared subprogram to be at the address that is specified for an unrelated subprogram (line 303).
- CD7105A This test requires that successive calls to CALENDAR.CLOCK change by at least SYSTEM.TICK; however, by Commentary AI-00201, it is only the expected frequency of change that must be at least SYSTEM.TICK--particular instances of change may be less (line 29).
- CD7203B, & CD7204B These tests use the 'SIZE length clause and attribute, whose interpretation is considered problematic by the WG9 ARG.
- CD7205D This test checks an invalid test objective: it treats the specification of storage to be reserved for a task's activation as though it were like the specification of storage for a collection.
- CE2107I This test requires that objects of two similar scalar types be distinguished when read from a file--DATA_ERROR is expected to be raised by an attempt to read one object as of the other type. However, it is not clear exactly how the Ada standard 14.2.4:4 is to be interpreted; thus, this test objective is not considered valid. (line 90).
- CE3111C This test requires certain behavior, when two files are associated with the same external file, that is not required by the Ada standard.
- CE3301A This test contains several calls to END_OF_LINE & END_OF_PAGE that have no parameter: these calls were intended to specify a file, not to refer to STANDARD_INPUT (lines 103, 107, 118, 132, & 136).
- CE3411B This test requires that a text file's column number be set to COUNT'LAST in order to check that LAYOUT_ERROR is raised by a subsequent PUT operation. But the former operation will generally raise an exception due to a lack of available disk space, and the test would thus encumber validation testing.

APPENDIX D

APPENDIX F OF THE Ada STANDARD

The only allowed implementation dependencies correspond to implementation-dependent pragmas, to certain machine-dependent conventions as mentioned in chapter 13 of the Ada Standard, and to certain allowed restrictions on representation clauses. The implementation-dependent characteristics of the AlsyCOMP_016, Version 4.1, are described in the following sections, which discuss topics in Appendix F of the Ada Standard. Implementation-specific portions of the package STANDARD are also included in this appendix.

package STANDARD is

...

type SHORT_SHORT_INTEGER is range -128 .. 127;

type SHORT_INTEGER is range -32_768 .. 32_767;

type INTEGER is range -2_147_483_648 .. 2_147_483_647;

type FLOAT is digits 6 range

-2#1.111_1111_1111_1111_1111_1111#E+127

..

2#1.111_1111_1111_1111_1111_1111#E+127;

type LONG_FLOAT is digits 15 range

-2#1.1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111#E1023

..

2#1.1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111#E1023;

type DURATION is delta 0.001 range -2097152.0 .. +2097152.0;

...

end STANDARD;

Copyright 1988 by Alsys

All rights reserved. No part of this document may be reproduced in any form or by any means without permission in writing from Alsys.

Printed: November 1988

Alsys reserves the right to make changes in specifications and other information contained in this publication without prior notice. Consult Alsys to determine whether such changes have been made.

Alsys, AdaWorld AdaProbe, AdaXref, AdaReformat, and AdaMake are registered trademarks of Alsys.

Unix is a registered trademark of AT&T.

386/ix is a registered trademark of Interactive Systems Corporation.

Microsoft is a registered trademark of Microsoft Corporation.

IBM, PC AT and PS/2 are registered trademarks of International Business Machines Corporation.

INTEL is a registered trademark of Intel Corporation.

TABLE OF CONTENTS

APPENDIX F	1
1 Implementation-Dependent Pragmas	2
1.1 INLINE	2
1.2 INTERFACE	2
1.3 INTERFACE_NAME	2
1.4 INDENT	3
1.5 Other Pragmas	4
2 Implementation-Dependent Attributes	4
3 Specification of the package SYSTEM	4
4 Restrictions on Representation Clauses	7
4.1 Enumeration Types	8
4.2 Integer Types	10
4.3 Floating Point Types	12
4.4 Fixed Point Types	13
4.5 Access Types	15
4.6 Task Types	16
4.7 Array Types	17
4.8 Record Types	20
4.8.1 RECORD_SIZE	25
4.8.2 VARIANT_INDEX	25
4.8.3 ARRAY_DESCRIPTOR	26
4.8.4 RECORD_DESCRIPTOR	27
5 Conventions for Implementation-Generated Names	28
6 Address Clauses	28
6.1 Address Clauses for Objects	28
6.2 Address Clauses for Program Units	29
6.3 Address Clauses for Entries	29

7	Restrictions on Unchecked Conversions	29
8	Input-Output Packages	29
8.1	Introduction	29
8.2	The FORM Parameter	30
9	Characteristics of Numeric Types	37
9.1	Integer Types	37
9.2	Floating Point Type Attributes	37
9.3	Attributes of Type DURATION	38
10	Other Implementation-Dependent Characteristics	38
10.1	Use of the Floating-Point Coprocessor (80287, 80387)	38
10.2	Characteristics of the Heap	39
10.3	Characteristics of Tasks	39
10.4	Definition of a Main Subprogram	40
10.5	Ordering of Compilation Units	40
11	Limitations	40
11.1	Compiler Limitations	40
11.2	Hardware Related Limitations	40
INDEX		41

APPENDIX F

Implementation - Dependent Characteristics

This appendix summarizes the implementation-dependent characteristics of the Alsys 386 Unix Ada Compiler. This appendix is a required part of the *Reference Manual for the Ada Programming Language* (called the *RM* in this appendix).

The sections of this appendix are as follows:

1. The form, allowed places, and effect of every implementation-dependent pragma.
2. The name and the type of every implementation-dependent attribute.
3. The specification of the package SYSTEM.
4. The list of all restrictions on representation clauses.
5. The conventions used for any implementation-generated name denoting implementation-dependent components.
6. The interpretation of expressions that appear in address clauses, including those for interrupts.
7. Any restrictions on unchecked conversions.
8. Any implementation-dependent characteristics of the input-output packages.
9. Characteristics of numeric types.
10. Other implementation-dependent characteristics.
11. Compiler limitations.

The name *Alsys Runtime Executive Programs* or simply *Runtime Executive* refers to the runtime library routines provided for all Ada programs. These routines implement the Ada heap, exceptions, tasking control, and other utility functions.

General systems programming notes are given in another document, the *Application Developer's Guide* (for example, parameter passing conventions needed for interface with assembly routines).

1 Implementation-Dependent Pragmas

1.1 INLINE

Pragma `INLINE` is fully supported; however, it is not possible to inline `INLINE` a subprogram in a declarative part.

1.2 INTERFACE

Ada programs can interface with subprograms written in Assembler and other languages through the use of the predefined pragma `INTERFACE` and the implementation-defined pragma `INTERFACE_NAME`.

Pragma `INTERFACE` specifies the name of an interfaced subprogram and the name of the programming language for which parameter passing conventions will be generated. Pragma `INTERFACE` takes the form specified in the RM:

```
pragma INTERFACE (language_name, subprogram_name);
```

where,

- *language_name* is `ASSEMBLER`, `ADA`, or `C`.
- *subprogram_name* is the name used within the Ada program to refer to the interfaced subprogram.

The only language names accepted by pragma `INTERFACE` are `ASSEMBLER`, `ADA` and `C`. The full implementation requirements for writing pragma `INTERFACE` subprograms are described in the *Application Developer's Guide*.

The language name used in the pragma `INTERFACE` does not have to have any relationship to the language actually used to write the interfaced subprogram. It is used only to tell the Compiler how to generate subprogram calls; that is, what kind of parameter passing techniques to use. The programmer can interface Ada programs with subroutines written in any other (compiled) language by understanding the mechanisms used for parameter passing by the Alsys 386 Unix Ada Compiler and the corresponding mechanisms of the chosen external language.

1.3 INTERFACE_NAME

Pragma `INTERFACE_NAME` associates the name of the interfaced subprogram with the external name of the interfaced subprogram. If pragma `INTERFACE_NAME` is not used, then the two names are assumed to be identical. This pragma takes the form:

```
pragma INTERFACE_NAME (subprogram_name, string_literal);
```

where,

- *subprogram_name* is the name used within the Ada program to refer to the interfaced subprogram.
- *string_literal* is the name by which the interfaced subprogram is referred to at link time.

The pragma `INTERFACE_NAME` is used to identify routines in other languages that are not named with legal Ada identifiers. Ada identifiers can only contain letters, digits, or underscores, whereas the Unix Linker allows external names to contain other characters, for example, the dollar sign (\$) or commercial at sign (@). These characters can be specified in the *string_literal* argument of the pragma `INTERFACE_NAME`.

The pragma `INTERFACE_NAME` is allowed at the same places of an Ada program as the pragma `INTERFACE`. (Location restrictions can be found in section 13.9 of the RM.) However, the pragma `INTERFACE_NAME` must always occur after the pragma `INTERFACE` declaration for the interfaced subprogram.

The *string_literal* of the pragma `INTERFACE_NAME` is passed through unchanged to the 386 Unix object file. (The 386 UNIX tools usually ignore the case of external identifiers. However recent versions of the Microsoft Linkers have options to treat external identifiers in a case-sensitive manner.) The maximum length of the *string_literal* is 40 characters. This limit is not checked by the Compiler, but the string is truncated by the Binder to meet the Intel object module format standard. Certain 386 UNIX tools have smaller limits.

The *Runtime Executive* contains several external identifiers. All such identifiers begin with either the string "ADA@" or the string "ADAS@". Accordingly, names prefixed by "ADA@" or "ADAS@" should be avoided by the user.

Example

```
package SAMPLE_DATA is
  function SAMPLE_DEVICE (X: INTEGER) return INTEGER;
  function PROCESS_SAMPLE (X: INTEGER) return INTEGER;
private
  pragma INTERFACE (ASSEMBLER, SAMPLE_DEVICE);
  pragma INTERFACE (ADA, PROCESS_SAMPLE);
  pragma INTERFACE_NAME (SAMPLE_DEVICE, "DEVIO$GET_SAMPLE");
end SAMPLE_DATA;
```

1.4 INDENT

Pragma `INDENT` is only used with *AdaReformat*. *AdaReformat* is the Alsys reformatter which offers the functionalities of a pretty-printer in an Ada environment.

The pragma is placed in the source file and interpreted by the Reformatter.

```
pragma INDENT(OFF);
```

causes *AdaReformat* not to modify the source lines after this pragma.

pragma INDENT(ON);

causes *AdaReformat* to resume its action after this pragma.

1.5 Other Pragmas

Pragmas IMPROVE and PACK are discussed in detail in the section on representation clauses and records (Chapter 5).

Pragma PRIORITY is accepted with the range of priorities running from 1 to 10 (see the definition of the predefined package SYSTEM in Section 3). Undefined priority (no pragma PRIORITY) is treated as though it were less than any defined priority value.

In addition to pragma SUPPRESS, it is possible to suppress all checks in a given compilation by the use of the Compiler option CHECKS.

2 Implementation-Dependent Attributes

P'IS_ARRAY For a prefix P that denotes any type or subtype, this attribute yields the value TRUE if P is an array type or an array subtype; otherwise, it yields the value FALSE.

P'RECORD_DESCRIPTOR These attributes are used to control the representation of implicit components of a record, see section 4.8
P'ARRAY_DESCRIPTOR

3 Specification of the package SYSTEM

The implementation does not allow the recompilation of package SYSTEM.

```
package SYSTEM is

  -- *****
  -- * (1) Required Definitions. *
  -- *****

  type NAME is (I_80x86);
  SYSTEM_NAME : constant NAME := I_80x86;

  STORAGE_UNIT : constant := 8;
  MEMORY_SIZE : constant := 640 * 1024;
```

```

-- System-Dependent Named Numbers:

MIN_INT      : constant := -(2**31);
MAX_INT      : constant := 2**31 - 1;
MAX_DIGITS   : constant := 15;
MAX_MANTISSA : constant := 31;
FINE_DELTA   : constant := 2#1.0#E-31;

-- For the high-resolution timer, the clock resolution is
-- 1.0 / 1024.0.
TICK         : constant := 1.0 / 18.2;

-- Other System-Dependent Declarations:

subtype PRIORITY is INTEGER range 1 .. 10;

-- The type ADDRESS is, in fact, implemented as a
-- 386 bit offset

type ADDRESS is private;
NULL_ADDRESS: constant ADDRESS := null;

-- *****
-- * (2) MACHINE TYPE CONVERSIONS *
-- *****

-- If the word / double-word operations below are used on
-- ADDRESS, then MSW yields the segment and LSW yields the
-- offset.

-- In the operations below, a BYTE_TYPE is any simple type
-- implemented on 8-bits (for example, SHORTSHORT_INTEGER), a WORD_TYPE is
-- any simple type implemented on 16-bits (for example, SHORT_INTEGER), and
-- a DOUBLE_WORD_TYPE is any simple type implemented on
-- 32-bits (for example, INTEGER, FLOAT, ADDRESS).

-- Byte <==> Word conversions:

-- Get the most significant byte:
generic
    type BYTE_TYPE is private;
    type WORD_TYPE is private;
function MSB (W: WORD_TYPE) return BYTE_TYPE;

-- Get the least significant byte:
generic
    type BYTE_TYPE is private;
    type WORD_TYPE is private;
function LSB (W: WORD_TYPE) return BYTE_TYPE;

```

```

-- Compose a word from two bytes:
generic
    type BYTE_TYPE is private;
    type WORD_TYPE is private;
function WORD (MSB, LSB: BYTE_TYPE) return WORD_TYPE;

-- Word <==> Double-Word conversions:

-- Get the most significant word:
generic
    type WORD_TYPE is private;
    type DOUBLE_WORD_TYPE is private;
function MSW (W: DOUBLE_WORD_TYPE) return WORD_TYPE;

-- Get the least significant word:
generic
    type WORD_TYPE is private;
    type DOUBLE_WORD_TYPE is private;
function LSW(W: DOUBLE_WORD_TYPE) return WORD_TYPE;

-- Compose a DATA double word from two words.
generic
    type WORD_TYPE is private;
    -- The following type must be a data type
    -- (for example, LONG_INTEGER):
    type DATA_DOUBLE_WORD is private;
function DOUBLE_WORD (MSW, LSW: WORD_TYPE)
    return DATA_DOUBLE_WORD;

-- Compose a REFERENCE double word from two words.
generic
    type WORD_TYPE is private;
    -- The following type must be a reference type
    -- (for example, access or ADDRESS):
    type REF_DOUBLE_WORD is private;
function REFERENCE (SEGMENT, OFFSET: WORD_TYPE)
    return REF_DOUBLE_WORD;

-- *****
-- * (3) OPERATIONS ON ADDRESS *
-- *****

-- You can get an address via 'ADDRESS attribute or by
-- instantiating the function REFERENCE, above, with
-- appropriate types.

-- Some addresses are used by the Compiler. For example,
-- the display is located at the low end of the DS segment,
-- and addresses SS:0 through SS:128 hold the task control

```

```

-- block and other information. Writing into these areas
-- will have unpredictable results.

-- Note that no operations are defined to get the values of
-- the segment registers, but if it is necessary an
-- interfaced function can be written.

generic
  type OBJECT is private;
  function FETCH_FROM_ADDRESS (FROM: ADDRESS) return OBJECT;

generic
  type OBJECT is private;
  procedure ASSIGN_TO_ADDRESS (OBJ: OBJECT; TO: ADDRESS);

private

  ...

end SYSTEM;

```

4 Restrictions on Representation Clauses

This section explains how objects are represented and allocated by the Alsys 386 Unix Ada compiler and how it is possible to control this using representation clauses.

The representation of an object is closely connected with its type. For this reason this section addresses successively the representation of enumeration, integer, floating point, fixed point, access, task, array and record types. For each class of type the representation of the corresponding objects is described.

Except in the case of array and record types, the description for each class of type is independent of the others. To understand the representation of an array type it is necessary to understand first the representation of its components. The same rule applies to record types.

Apart from implementation defined pragmas, Ada provides three means to control the size of objects:

- a (predefined) pragma PACK, applicable to array types
- a record representation clause
- a size specification

For each class of types the effect of a size specification is described. Interference between size specifications, packing and record representation clauses is described under array and record types.

Representation clauses on derived record, or derived tasks are not supported.

Size representation clauses on types derived from private types are not supported when the derived type is declared outside the private part of the defining package.

4.1 Enumeration Types

Internal codes of enumeration literals

When no enumeration representation clause applies to an enumeration type, the internal code associated with an enumeration literal is the position number of the enumeration literal. Then, for an enumeration type with n elements, the internal codes are the integers 0, 1, 2, .., $n-1$.

An enumeration representation clause can be provided to specify the value of each internal code as described in RM 13.3. The Alsys compiler fully implements enumeration representation clauses.

As internal codes must be machine integers the internal codes provided by an enumeration representation clause must be in the range $-2^{31} .. 2^{31}-1$.

Encoding of enumeration values

An enumeration value is always represented by its internal code in the program generated by the compiler.

Minimum size of an enumeration subtype

The minimum size of an enumeration subtype is the minimum number of bits that is necessary for representing the internal codes of the subtype values in normal binary form.

For a static subtype, if it has a null range its minimum size is 1. Otherwise, if m and M are the values of the internal codes associated with the first and last enumeration values of the subtype, then its minimum size L is determined as follows. For $m \geq 0$, L is the smallest positive integer such that $M \leq 2^L - 1$. For $m < 0$, L is the smallest positive integer such that $-2^{L-1} \leq m$ and $M \leq 2^{L-1} - 1$.

```
type COLOR is (GREEN, BLACK, WHITE, RED, BLUE, YELLOW);  
-- The minimum size of COLOR is 3 bits.
```

```
subtype BLACK_AND_WHITE is COLOR range BLACK .. WHITE;  
-- The minimum size of BLACK_AND_WHITE is 2 bits.
```

```
subtype BLACK_OR_WHITE is BLACK_AND_WHITE range X .. X;  
-- Assuming that X is not static, the minimum size of BLACK_OR_WHITE is  
-- 2 bits (the same as the minimum size of its type mark  
BLACK_AND_WHITE).
```

Size of an enumeration subtype

When no size specification is applied to an enumeration type or first named subtype, the objects of that type or first named subtype are represented as signed machine integers. The machine provides 8, 16 and 32 bit integers, and the compiler selects automatically the smallest signed machine integer which can hold each of the internal codes of the enumeration type (or subtype). The size of the enumeration type and of any of its subtypes is thus 8, 16 or 32 bits.

When a size specification is applied to an enumeration type, this enumeration type and each of its subtypes has the size specified by the length clause. The same rule applies to a first named subtype. The size specification must of course specify a value greater than or equal to the minimum size of the type or subtype to which it applies:

type EXTENDED is

(-- The usual ASCII characters.

NUL,	SOH,	STX,	ETX,	EOT,	ENQ,	ACK,	BEL,
BS,	HT,	LF,	VT,	FF,	CR,	SO,	SI,
DLE,	DC1,	DC2,	DC3,	DC4,	NAK,	SYN,	ETB,
CAN,	EM,	SUB,	ESC,	FS,	GS,	RS,	US,
' ',	'!',	'"',	'#',	'\$',	'%',	'&',	'"',
'(',	')',	'*',	'+',	'-',	'.',	'/',	'/',
'0',	'1',	'2',	'3',	'4',	'5',	'6',	'7',
'8',	'9',	'<',	'>',	'?',	'@',	'A',	'B',
'C',	'D',	'E',	'F',	'G',	'H',	'I',	'J',
'K',	'L',	'M',	'N',	'O',	'P',	'Q',	'R',
'S',	'T',	'U',	'V',	'W',	'X',	'Y',	'Z',
'[',	'\',	']',	'^',	'_',	'a',	'b',	'c',
'd',	'e',	'f',	'g',	'h',	'i',	'j',	'k',
'l',	'm',	'n',	'o',	'p',	'q',	'r',	's',
't',	'u',	'v',	'w',	'x',	'y',	'z',	'{',
' ',	'}',	'~',	DEL,				

-- Extended characters

LEFT_ARROW,
RIGHT_ARROW,
UPPER_ARROW,
LOWER_ARROW,
UPPER_LEFT_CORNER,
UPPER_RIGHT_CORNER,
LOWER_RIGHT_CORNER,
LOWER_LEFT_CORNER

);

for EXTENDED'SIZE use 8;

-- The size of type EXTENDED will be one byte. Its objects will be represented
-- as unsigned 8 bit integers.

The Alsys compiler fully implements size specifications. Nevertheless, as enumeration values are coded using integers, the specified length cannot be greater than 32 bits.

Size of the objects of an enumeration subtype

Provided its size is not constrained by a record component clause or a pragma PACK, an object of an enumeration subtype has the same size as its subtype.

4.2 Integer Types

Predefined integer types

There are three predefined integer types in the Alsys implementation for I80x86 machines:

```
type SHORT_SHORT_INTEGER is range -2**07 .. 2**07-1;  
type SHORT_INTEGER       is range -2**15 .. 2**15-1;  
type INTEGER              is range -2**31 .. 2**31-1;
```

Selection of the parent of an integer type

An integer type declared by a declaration of the form:

```
type T is range L .. R;
```

is implicitly derived from a predefined integer type. The compiler automatically selects the predefined integer type whose range is the smallest that contains the values L to R inclusive.

Encoding of integer values

Binary code is used to represent integer values. Negative numbers are represented using two's complement.

Minimum size of an integer subtype

The minimum size of an integer subtype is the minimum number of bits that is necessary for representing the internal codes of the subtype values in normal binary form.

For a static subtype, if it has a null range its minimum size is 1. Otherwise, if m and M are the lower and upper bounds of the subtype, then its minimum size L is determined as follows. For $m \geq 0$, L is the smallest positive integer such that $M \leq 2^L - 1$. For $m < 0$, L is the smallest positive integer that $-2^{L-1} \leq m$ and $M \leq 2^{L-1} - 1$.

subtype S is INTEGER range 0 .. 7;
-- The minimum size of S is 3 bits.

subtype D is S range X .. Y;
-- Assuming that X and Y are not static, the minimum size of
-- D is 3 bits (the same as the minimum size of its type mark S).

Size of an integer subtype

The sizes of the predefined integer types `SHORT_SHORT_INTEGER`, `SHORT_INTEGER` and `INTEGER` are respectively 8, 16 and 32 bits.

When no size specification is applied to an integer type or to its first named subtype (if any), its size and the size of any of its subtypes is the size of the predefined type from which it derives, directly or indirectly. For example:

type S is range 80 .. 100;
-- S is derived from `SHORT_SHORT_INTEGER`, its size is 8 bits.

type J is range 0 .. 255;
-- J is derived from `SHORT_INTEGER`, its size is 16 bits.

type N is new J range 80 .. 100;
-- N is indirectly derived from `SHORT_INTEGER`, its size is 16 bits.

When a size specification is applied to an integer type, this integer type and each of its subtypes has the size specified by the length clause. The same rule applies to a first named subtype. The size specification must of course specify a value greater than or equal to the minimum size of the type or subtype to which it applies:

type S is range 80 .. 100;
for S'SIZE use 32;
-- S is derived from `SHORT_INTEGER`, but its size is 32 bits
-- because of the size specification.

type J is range 0 .. 255;
for J'SIZE use 8;
-- J is derived from `INTEGER`, but its size is 8 bits because
-- of the size specification.

type N is new J range 80 .. 100;
-- N is indirectly derived from `INTEGER`, but its size is 8 bits
-- because N inherits the size specification of J.

Size of the objects of an integer subtype

Provided its size is not constrained by a record component clause or a pragma `PACK`, an object of an integer subtype has the same size as its subtype.

4.3 Floating Point Types

Predefined floating point types

There are two predefined floating point types in the Alsys implementation for I80x86 machines:

```
type FLOAT is
  digits 6 range -(2.0 - 2.0**(-23))*2.0**127 .. (2.0 - 2.0**(-23))*2.0**127;

type LONG_FLOAT is
  digits 15 range -(2.0 - 2.0**(-51))*2.0**1023 .. (2.0 - 2.0**(-51))*2.0**1023;
```

Selection of the parent of a floating point type

A floating point type declared by a declaration of the form:

```
type T is digits D [range L .. R];
```

is implicitly derived from a predefined floating point type. The compiler automatically selects the smallest predefined floating point type whose number of digits is greater than or equal to D and which contains the values L to R inclusive.

Encoding of floating point values

In the program generated by the compiler, floating point values are represented using the IEEE standard formats for single and double floats.

The values of the predefined type FLOAT are represented using the single float format. The values of the predefined type LONG_FLOAT are represented using the double float format. The values of any other floating point type are represented in the same way as the values of the predefined type from which it derives, directly or indirectly.

Minimum size of a floating point subtype

The minimum size of a floating point subtype is 32 bits if its base type is FLOAT or a type derived from FLOAT; it is 64 bits if its base type is LONG_FLOAT or a type derived from LONG_FLOAT.

Size of a floating point subtype

The sizes of the predefined floating point types FLOAT and LONG_FLOAT are respectively 32 and 64 bits.

The size of a floating point type and the size of any of its subtypes is the size of the predefined type from which it derives directly or indirectly.

The only size that can be specified for a floating point type or first named subtype using a size specification is its usual size (32 or 64 bits).

Size of the objects of a floating point subtype

An object of a floating point subtype has the same size as its subtype.

4.4 Fixed Point Types

Small of a fixed point type

If no specification of small applies to a fixed point type, then the value of small is determined by the value of delta as defined by RM 3.5.9.

A specification of small can be used to impose a value of small. The value of small is required to be a power of two.

Predefined fixed point types

To implement fixed point types, the Alsys compiler for I80x86 machines uses a set of anonymous predefined types of the form:

```
type SHORT_FIXED is delta D range  $(-2.0^{**07}-1)*S$  ..  $2.0^{**07}*S$ ;  
for SHORT_FIXED'SMALL use S;
```

```
type FIXED is delta D range  $(-2.0^{**15}-1)*S$  ..  $2.0^{**15}*S$ ;  
for FIXED'SMALL use S;
```

```
type LONG_FIXED is delta D range  $(-2.0^{**31}-1)*S$  ..  $2.0^{**31}*S$ ;  
for LONG_FIXED'SMALL use S;
```

where D is any real value and S any power of two less than or equal to D.

Selection of the parent of a fixed point type

A fixed point type declared by a declaration of the form:

```
type T is delta D range L .. R;
```

possibly with a small specification:

```
for T'SMALL use S;
```

is implicitly derived from a predefined fixed point type. The compiler automatically selects the predefined fixed point type whose small and delta are the same as the small and delta of T and whose range is the shortest that includes the values L to R inclusive.

Encoding of fixed point values

In the program generated by the compiler, a safe value V of a fixed point subtype F is represented as the integer:

$$V / F\text{'BASE'SMALL}$$

Minimum size of a fixed point subtype

The minimum size of a fixed point subtype is the minimum number of binary digits that is necessary for representing the values of the range of the subtype using the small of the base type.

For a static subtype, if it has a null range its minimum size is 1. Otherwise, s and S being the bounds of the subtype, if i and I are the integer representations of m and M , the smallest and the greatest model numbers of the base type such that $s < m$ and $M < S$, then the minimum size L is determined as follows. For $i \geq 0$, L is the smallest positive integer such that $I \leq 2^L - 1$. For $i < 0$, L is the smallest positive integer such that $-2^{L-1} \leq i$ and $I \leq 2^{L-1} - 1$.

```
type F is delta 2.0 range 0.0 .. 500.0;  
-- The minimum size of F is 8 bits.
```

```
subtype S is F delta 16.0 range 0.0 .. 250.0;  
-- The minimum size of S is 7 bits.
```

```
subtype D is S range X .. Y;  
-- Assuming that X and Y are not static, the minimum size of D is 7 bits  
-- (the same as the minimum size of its type mark S).
```

Size of a fixed point subtype

The sizes of the predefined fixed point types *SHORT_FIXED*, *FIXED* and *LONG_FIXED* are respectively 8, 16 and 32 bits.

When no size specification is applied to a fixed point type or to its first named subtype, its size and the size of any of its subtypes is the size of the predefined type from which it derives directly or indirectly. For example:

```
type S is delta 0.01 range 0.8 .. 1.0;  
-- S is derived from an 8 bit predefined fixed type, its size is 8 bits.
```

```
type F is delta 0.01 range 0.0 .. 2.0;  
-- F is derived from a 16 bit predefined fixed type, its size is 16 bits.
```

```
type N is new F range 0.8 .. 1.0;  
-- N is indirectly derived from a 16 bit predefined fixed type, its size is 16 bits.
```

When a size specification is applied to a fixed point type, this fixed point type and each of its subtypes has the size specified by the length clause. The same rule applies to a first named subtype. The size specification must of course specify a value greater than or equal to the minimum size of the type or subtype to which it applies:

```
type S is delta 0.01 range 0.8 .. 1.0;
for S'SIZE use 32;
-- S is derived from an 8 bit predefined fixed type, but its size is 32 bits
-- because of the size specification.
```

```
type F is delta 0.01 range 0.0 .. 2.0;
for F'SIZE use 8;
-- F is derived from a 16 bit predefined fixed type, but its size is 8 bits
-- because of the size specification.
```

```
type N is new F range 0.8 .. 1.0;
-- N is indirectly derived from a 16 bit predefined fixed type, but its size is
-- 8 bits because N inherits the size specification of F.
```

The Alsys compiler fully implements size specifications. Nevertheless, as fixed point objects are represented using machine integers, the specified length cannot be greater than 32 bits.

Size of the objects of a fixed point subtype

Provided its size is not constrained by a record component clause or a pragma PACK, an object of a fixed point type has the same size as its subtype.

4.5 Access Types

Collection Size

As described in RM 13.2, a specification of collection size can be provided in order to reserve storage space for the collection of an access type.

When no STORAGE_SIZE specification applies to an access type, no storage space is reserved for its collection, and the value of the attribute STORAGE_SIZE is then 0.

STORAGE_SIZE clause on collections of unconstrained objects is not supported by the implementation.

Encoding of access values.

Access values are machine addresses.

Minimum size of an access subtype

The minimum size of an access subtype is 32 bits.

Size of an access subtype

The size of an access subtype is 32 bits, the same as its minimum size.

The only size that can be specified for an access type using a size specification is its usual size (32 bits).

Size of an object of an access subtype

An object of an access subtype has the same size as its subtype, thus an object of an access subtype is always 32 bits long.

4.6 Task Types

Storage for a task activation

When no length clause is used to specify the storage space to be reserved for a task activation, the storage space indicated at bind time is used for this activation.

As described in RM 13.2, a length clause can be used to specify the storage space for the activation of each of the tasks of a given type. In this case the value indicated at bind time is ignored for this task type, and the length clause is obeyed.

It is not allowed to apply such a length clause to a derived type. The same storage space is reserved for the activation of a task of a derived type as for the activation of a task of the parent type.

Encoding of task values

Encoding of a task value is not described here.

Minimum size of a task subtype

The minimum size of a task subtype is 32 bits.

Size of a task subtype

The size of a task subtype is 32 bits, the same as its minimum size.

A size specification has no effect on a task type. The only size that can be specified using such a length clause is its minimum size.

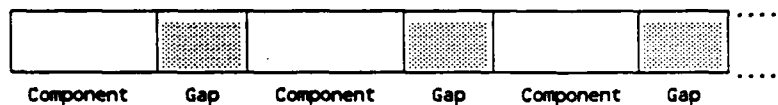
Size of the objects of a task subtype

An object of a task subtype has the same size as its subtype. Thus an object of a task subtype is always 32 bits long.

4.7 Array Types

Layout of an array

Each array is allocated in a contiguous area of storage units. All the components have the same size. A gap may exist between two consecutive components (and after the last one). All the gaps have the same size.



Components

If the array is not packed, the size of the components is the size of the subtype of the components:

```
type A is array (1 .. 8) of BOOLEAN;
-- The size of the components of A is the size of the type BOOLEAN: 8 bits.

type DECIMAL_DIGIT is range 0 .. 9;
for DECIMAL_DIGIT'SIZE use 4;
type BINARY_CODED_DECIMAL is
  array (INTEGER range <>) of DECIMAL_DIGIT;
-- The size of the type DECIMAL_DIGIT is 4 bits. Thus in an array of
-- type BINARY_CODED_DECIMAL each component will be represented on
-- 4 bits as in the usual BCD representation.
```

If the array is packed and its components are neither records nor arrays, the size of the components is the minimum size of the subtype of the components:

```
type A is array (1 .. 8) of BOOLEAN;
pragma PACK(A);
-- The size of the components of A is the minimum size of the type BOOLEAN:
-- 1 bit.

type DECIMAL_DIGIT is range 0 .. 9;
for DECIMAL_DIGIT'SIZE use 32;
type BINARY_CODED_DECIMAL is
  array (INTEGER range <>) of DECIMAL_DIGIT;
pragma PACK(BINARY_CODED_DECIMAL);
-- The size of the type DECIMAL_DIGIT is 32 bits, but, as
```

```
-- BINARY_CODED_DECIMAL is packed, each component of an array of this
-- type will be represented on 4 bits as in the usual BCD representation.
```

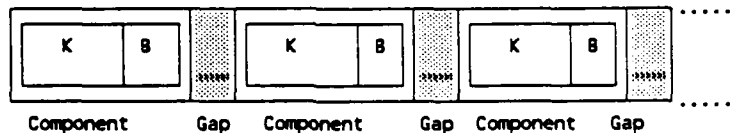
Packing the array has no effect on the size of the components when the components are records or arrays.

Gaps

If the components are records or arrays, no size specification applies to the subtype of the components and the array is not packed, then the compiler may choose a representation with a gap after each component; the aim of the insertion of such gaps is to optimize access to the array components and to their subcomponents. The size of the gap is chosen so that the relative displacement of consecutive components is a multiple of the alignment of the subtype of the components. This strategy allows each component and subcomponent to have an address consistent with the alignment of its subtype:

```
type R is
  record
    K : SHORT_INTEGER; -- SHORT_INTEGER is even byte aligned.
    B : BOOLEAN; -- BOOLEAN is byte aligned.
  end record;
-- Record type R is even byte aligned. Its size is 24 bits.

type A is array (1 .. 10) of R;
-- A gap of one byte is inserted after each component in order to respect the
-- alignment of type R. The size of an array of type A will be 320 bits.
```



Array of type A: each subcomponent K has an even offset.

If a size specification applies to the subtype of the components or if the array is packed, no gaps are inserted:

```
type R is
  record
    K : SHORT_INTEGER;
    B : BOOLEAN;
  end record;

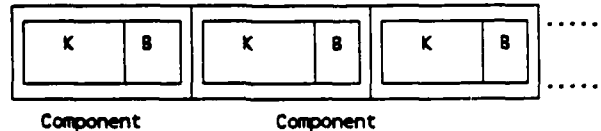
type A is array (1 .. 10) of R;
pragma PACK(A);
-- There is no gap in an array of type A because
-- A is packed.
```

-- The size of an object of type A will be 240 bits.

type NR is new R;
for NR'SIZE use 24;

type B is array (1 .. 10) of NR;

-- There is no gap in an array of type B because
-- NR has a size specification.
-- The size of an object of type B will be 240 bits.



Array of type A or B: a subcomponent K can have an odd offset.

Size of an array subtype

The size of an array subtype is obtained by multiplying the number of its components by the sum of the size of the components and the size of the gaps (if any). If the subtype is unconstrained, the maximum number of components is considered.

The size of an array subtype cannot be computed at compile time

- if it has non-static constraints or is an unconstrained array type with non-static index subtypes (because the number of components can then only be determined at run time).
- if the components are records or arrays and their constraints or the constraints of their subcomponents (if any) are not static (because the size of the components and the size of the gaps can then only be determined at run time).

As has been indicated above, the effect of a pragma PACK on an array type is to suppress the gaps and to reduce the size of the components. The consequence of packing an array type is thus to reduce its size.

If the components of an array are records or arrays and their constraints or the constraints of their subcomponents (if any) are not static, the compiler ignores any pragma PACK applied to the array type but issues a warning message. Apart from this limitation, array packing is fully implemented by the Alsys compiler.

A size specification applied to an array type or first named subtype has no effect. The only size that can be specified using such a length clause is its usual size. Nevertheless, such a length clause can be useful to verify that the layout of an array is as expected by the application.

Size of the objects of an array subtype

The size of an object of an array subtype is always equal to the size of the subtype of the object.

4.8 Record Types

Layout of a record

Each record is allocated in a contiguous area of storage units. The size of a record component depends on its type. Gaps may exist between some components.

The positions and the sizes of the components of a record type object can be controlled using a record representation clause as described in RM 13.4. In the Alsys implementation for 180x86 machines there is no restriction on the position that can be specified for a component of a record. If a component is not a record or an array, its size can be any size from the minimum size to the size of its subtype. If a component is a record or an array, its size must be the size of its subtype:

```
type INTERRUPT_MASK is array (0 .. 2) of BOOLEAN;
pragma PACK(INTERRUPT_MASK);
-- The size of INTERRUPT_MASK is 3 bits.

type CONDITION_CODE is 0 .. 1;
-- The size of CONDITION_CODE is 8 bits, its minimum size is 1 bit.

type STATUS_BIT is new BOOLEAN;
for STATUS_BIT'SIZE use 1;
-- The size and the minimum size of STATUS_BIT are 1 bit.

SYSTEM : constant := 0;
USER    : constant := 1;

type STATUS_REGISTER is
  record
    T : STATUS_BIT;           -- Trace
    S : STATUS_BIT;           -- Supervisor
    I : INTERRUPT_MASK;       -- Interrupt mask
    X : CONDITION_CODE;       -- Extend
    N : CONDITION_CODE;       -- Negative
    Z : CONDITION_CODE;       -- Zero
    V : CONDITION_CODE;       -- Overflow
    C : CONDITION_CODE;       -- Carry
  end record;

for STATUS_REGISTER use
  record at mod 2;
    T at SYSTEM range 0 .. 0;
    S at SYSTEM range 2 .. 2;
    I at SYSTEM range 5 .. 7;
```

```

X  at USER      range 3 .. 3;
N  at USER      range 4 .. 4;
Z  at USER      range 5 .. 5;
V  at USER      range 6 .. 6;
C  at USER      range 7 .. 7;
end record;

```

A record representation clause need not specify the position and the size for every component.

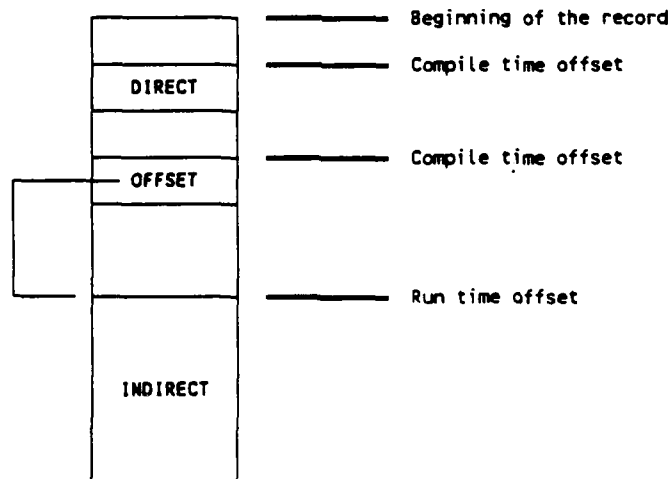
Pragma PACK has no effect on records.

If no component clause applies to a component of a record, its size is the size of its subtype. Its position is chosen by the compiler so as to optimize access to the components of the record: the offset of the component is chosen as a multiple of 8 bits if the objects of the component subtype are usually byte aligned, but a multiple of 16 bits if these objects are usually even byte aligned. Moreover, the compiler chooses the position of the component so as to reduce the number of gaps and thus the size of the record objects.

Because of these optimizations, there is no connection between the order of the components in a record type declaration and the positions chosen by the compiler for the components in a record object.

Indirect components

If the offset of a component cannot be computed at compile time, this offset is stored in the record objects at run time and used to access the component. Such a component is said to be indirect while other components are said to be direct:



A direct and an indirect component

If a record component is a record or an array, the size of its subtype may be evaluated at run time and may even depend on the discriminants of the record. We will call these components dynamic components:

```
type DEVICE is (SCREEN, PRINTER);

type COLOR is (GREEN, RED, BLUE);

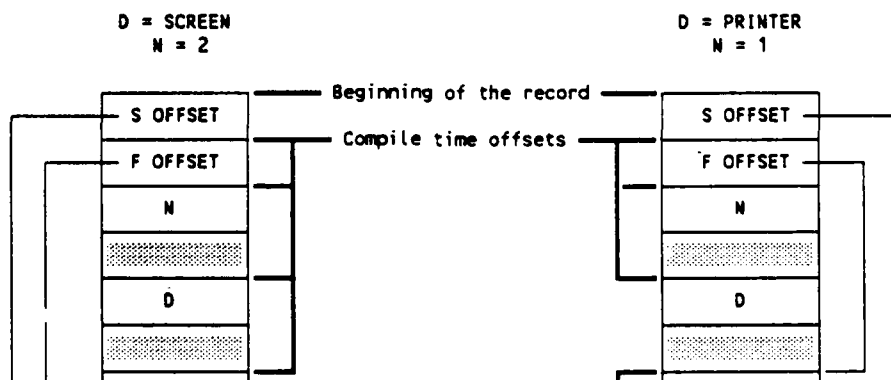
type SERIES is array (POSITIVE range <>) of INTEGER;

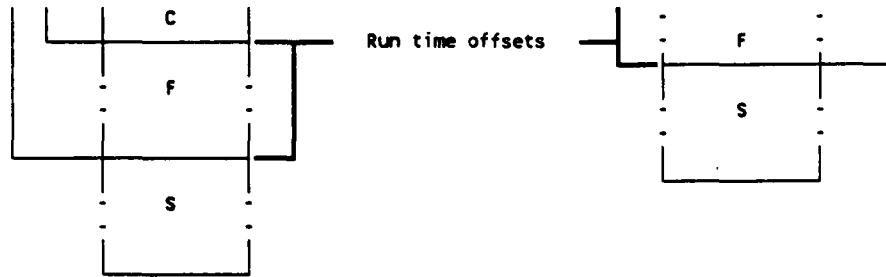
type GRAPH (L : NATURAL) is
  record
    X : SERIES(1 .. L); -- The size of X depends on L
    Y : SERIES(1 .. L); -- The size of Y depends on L
  end record;

Q : POSITIVE;

type PICTURE (N : NATURAL; D : DEVICE) is
  record
    F : GRAPH(N); -- The size of F depends on N
    S : GRAPH(Q); -- The size of S depends on Q
    case D is
      when SCREEN =>
        C : COLOR;
      when PRINTER =>
        null;
    end case;
  end record;
```

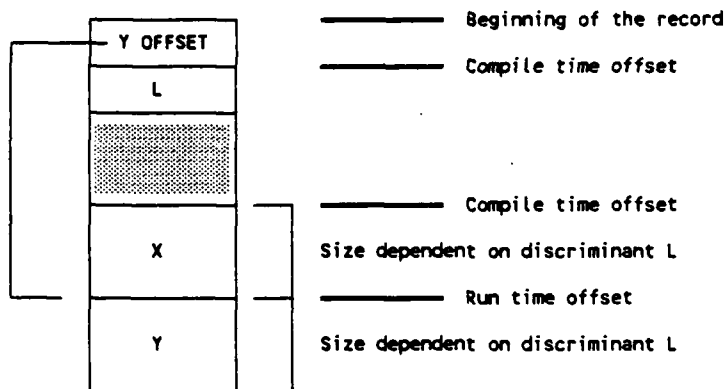
Any component placed after a dynamic component has an offset which cannot be evaluated at compile time and is thus indirect. In order to minimize the number of indirect components, the compiler groups the dynamic components together and places them at the end of the record:





The record type PICTURE: F and S are placed at the end of the record

Because of this approach, the only indirect components are dynamic components. But not all dynamic components are necessarily indirect: if there are dynamic components in a component list which is not followed by a variant part, then exactly one dynamic component of this list is a direct component because its offset can be computed at compilation time (the only dynamic components that are direct components are in this situation):



The record type GRAPH: the dynamic component X is a direct component.

The offset of an indirect component is always expressed in storage units.

The space reserved for the offset of an indirect component must be large enough to store the size of any value of the record type (the maximum potential offset). The compiler evaluates an upper bound MS of this size and treats an offset as a component having an anonymous integer type whose range is 0 .. MS.

If C is the name of an indirect component, then the offset of this component can be denoted in a component clause by the implementation generated name C'OFFSET.

Implicit components

In some circumstances, access to an object of a record type or to its components involves computing information which only depends on the discriminant values. To avoid useless recomputation the compiler stores this information in the record objects, updates it when the values of the discriminants are modified and uses it when the objects or its components are accessed. This information is stored in special components called implicit components.

An implicit component may contain information which is used when the record object or several of its components are accessed. In this case the component will be included in any record object (the implicit component is considered to be declared before any variant part in the record type declaration). There can be two components of this kind; one is called RECORD_SIZE and the other VARIANT_INDEX.

On the other hand an implicit component may be used to access a given record component. In that case the implicit component exists whenever the record component exists (the implicit component is considered to be declared at the same place as the record component). Components of this kind are called ARRAY_DESCRIPTORs or RECORD_DESCRIPTORs.

4.8.1 RECORD_SIZE

This implicit component is created by the compiler when the record type has a variant part and its discriminants are defaulted. It contains the size of the storage space necessary to store the current value of the record object (note that the storage effectively allocated for the record object may be more than this).

The value of a RECORD_SIZE component may denote a number of bits or a number of storage units. In general it denotes a number of storage units, but if any component clause specifies that a component of the record type has an offset or a size which cannot be expressed using storage units, then the value designates a number of bits.

The implicit component RECORD_SIZE must be large enough to store the maximum size of any value of the record type. The compiler evaluates an upper bound MS of this size and then considers the implicit component as having an anonymous integer type whose range is 0 .. MS.

If R is the name of the record type, this implicit component can be denoted in a component clause by the implementation generated name R'RECORD_SIZE.

4.8.2 VARIANT_INDEX

This implicit component is created by the compiler when the record type has a variant part. It indicates the set of components that are present in a record value. It is used when a discriminant check is to be done.

Component lists that do not contain a variant part are numbered. These numbers are the possible values of the implicit component VARIANT_INDEX.

```
type VEHICLE is (AIRCRAFT, ROCKET, BOAT, CAR);
```

```
type DESCRIPTION (KIND : VEHICLE := CAR) is
```

```
  record
```

```
    SPEED : INTEGER;
```

```
    case KIND is
```

```
      when AIRCRAFT | CAR =>
```

```
        WHEELS : INTEGER;
```

```
        case KIND is
```

```
          when AIRCRAFT => -- 1
```

```
            WINGSPAN : INTEGER;
```

```
          when others => -- 2
```

```
            null;
```

```
        end case;
```

```
      when BOAT => -- 3
```

```
        STEAM : BOOLEAN;
```

```
      when ROCKET => -- 4
```

```
        STAGES : INTEGER;
```

```
    end case;
```

```
  end record;
```

The value of the variant index indicates the set of components that are present in a record value:

Variant Index	Set
1	(KIND, SPEED, WHEELS, WINGSPAN)
2	(KIND, SPEED, WHEELS)
3	(KIND, SPEED, STEAM)
4	(KIND, SPEED, STAGES)

A comparison between the variant index of a record value and the bounds of an interval is enough to check that a given component is present in the value:

Component	Interval
KIND	--
SPEED	--
WHEELS	1 .. 2
WINGSPAN	1 .. 1
STEAM	3 .. 3
STAGES	4 .. 4

The implicit component `VARIANT_INDEX` must be large enough to store the number `V` of component lists that don't contain variant parts. The compiler treats this implicit component as having an anonymous integer type whose range is `1 .. V`.

If `R` is the name of the record type, this implicit component can be denoted in a component clause by the implementation generated name `R'VARIANT_INDEX`.

4.8.3 ARRAY_DESCRIPTOR

An implicit component of this kind is associated by the compiler with each record component whose subtype is an anonymous array subtype that depends on a discriminant of the record. It contains information about the component subtype.

The structure of an implicit component of kind `ARRAY_DESCRIPTOR` is not described in this documentation. Nevertheless, if a programmer is interested in specifying the location of a component of this kind using a component clause, he can obtain the size of the component using the `ASSEMBLY` parameter in the `COMPILE` command.

The compiler treats an implicit component of the kind `ARRAY_DESCRIPTOR` as having an anonymous array type. If `C` is the name of the record component whose subtype is described by the array descriptor, then this implicit component can be denoted in a component clause by the implementation generated name `C'ARRAY_DESCRIPTOR`.

4.8.4 RECORD_DESCRIPTOR

An implicit component of this kind is associated by the compiler with each record component whose subtype is an anonymous record subtype that depends on a discriminant of the record. It contains information about the component subtype.

The structure of an implicit component of kind RECORD_DESCRIPTOR is not described in this documentation. Nevertheless, if a programmer is interested in specifying the location of a component of this kind using a component clause, he can obtain the size of the component using the ASSEMBLY parameter in the COMPILE command.

The compiler treats an implicit component of the kind RECORD_DESCRIPTOR as having an anonymous array type. If C is the name of the record component whose subtype is described by the record descriptor, then this implicit component can be denoted in a component clause by the implementation generated name C'RECORD_DESCRIPTOR.

Suppression of implicit components

The Alsys implementation provides the capability of suppressing the implicit components RECORD_SIZE and/or VARIANT_INDEX from a record type. This can be done using an implementation defined pragma called IMPROVE. The syntax of this pragma is as follows:

```
pragma IMPROVE ( TIME | SPACE , [ON =>] simple_name );
```

The first argument specifies whether TIME or SPACE is the primary criterion for the choice of the representation of the record type that is denoted by the second argument.

If TIME is specified, the compiler inserts implicit components as described above. If on the other hand SPACE is specified, the compiler only inserts a VARIANT_INDEX or a RECORD_SIZE.

component if this component appears in a record representation clause that applies to the record type. A record representation clause can thus be used to keep one implicit component while suppressing the other.

A pragma IMPROVE that applies to a given record type can occur anywhere that a representation clause is allowed for this type.

Size of a record subtype

Unless a component clause specifies that a component of a record type has an offset or a size which cannot be expressed using storage units, the size of a record subtype is rounded up to a whole number of storage units.

The size of a constrained record subtype is obtained by adding the sizes of its components and the sizes of its gaps (if any). This size is not computed at compile time

- when the record subtype has non-static constraints,
- when a component is an array or a record and its size is not computed at compile time.

The size of an unconstrained record subtype is obtained by adding the sizes of the components and the sizes of the gaps (if any) of its largest variant. If the size of a component or of a gap cannot be evaluated exactly at compile time an upper bound of this size is used by the compiler to compute the subtype size.

A size specification applied to a record type or first named subtype has no effect. The only size that can be specified using such a length clause is its usual size. Nevertheless, such a length clause can be useful to verify that the layout of a record is as expected by the application.

Size of an object of a record subtype

An object of a constrained record subtype has the same size as its subtype.

An object of an unconstrained record subtype has the same size as its subtype if this size is less than or equal to 8 kb. If the size of the subtype is greater than this, the object has the size necessary to store its current value; storage space is allocated and released as the discriminants of the record change.

5 Conventions for Implementation-Generated Names

The Alsys 386 Unix Ada Compiler may add fields to record objects and have descriptors in memory for record or array objects. These fields are not accessible to the user through any implementation-generated name or attribute.

The following predefined packages are reserved to Alsys and cannot be recompiled in Version 4.2:

```

ALSYS_ADA_RUNTIME
ALSYS_BASIC_IO
ALSYS_BASIC_DIRECT_IO
ALSYS_BASIC_SEQUENTIAL_IO

```

6 Address Clauses

6.1 Address Clauses for Objects

An address clause can be used to specify an address for an object as described in RM 13.5. When such a clause applies to an object no storage is allocated for it in the program generated by the compiler. The program accesses the object using the address specified in the clause.

An address clause is not allowed for task objects, for unconstrained records whose size is greater than 8 kb., or for a constant.

6.2 Address Clauses for Program Units

Address clauses for program units are not implemented in the current version of the compiler.

6.3 Address Clauses for Entries

Address clauses for entries are not implemented in the current version of the compiler.

7 Restrictions on Unchecked Conversions

Unchecked conversions are allowed between any types. It is the programmer's responsibility to determine if the desired effect is achieved.

8 Input-Output Packages

In this part of the Appendix the implementation-specific aspects of the input-output system are described.

8.1 Introduction

In Ada, input-output operations (IO) are considered to be performed on *objects* of a certain file type rather than being performed directly on external files. An external file is anything external to the program that can produce a value to be read or receive a value to be written. Values transferred for a given file must be all of one type.

Generally, in Ada documentation, the term *file* refers to an object of a certain file type, whereas a physical manifestation is known as an *external file*. An external file is characterized by

- Its NAME, which is a string defining a legal path name under the current version of the operating system.
- Its FORM, which gives implementation-dependent information on file characteristics.

Both the NAME and the FORM appear explicitly as parameters of the Ada CREATE and OPEN procedures. Though a file is an object of a certain file type, ultimately the object has to correspond to an external file. Both CREATE and OPEN associate a NAME of an external file (of a certain FORM) with a program file object.

Ada IO operations are provided by means of standard packages [14].

SEQUENTIAL_IO	A generic package for sequential files of a single element type.
DIRECT_IO	A generic package for direct (random) access files.
TEXT_IO	A generic package for human readable (text, ASCII) files.
IO_EXCEPTIONS	A package which defines the exceptions needed by the above three packages.

The generic package **LOW_LEVEL_IO** is not implemented in this version.

The upper bound for index values in **DIRECT_IO** and for line, column and page numbers in **TEXT_IO** is given by

COUNT'LAST = 231 - 1**

The upper bound for field widths in **TEXT_IO** is given by

FIELD'LAST = 255

8.2 The FORM Parameter

The **FORM** parameter of both the **CREATE** and **OPEN** procedures in Ada specifies the characteristics of the external file involved.

The **CREATE** procedure establishes a new external file, of a given **NAME** and **FORM**, and associates it with a specified program file object. The external file is created (and the file object set) with a specified (or default) file mode. If the external file already exists, the file will be erased. The exception **USE_ERROR** is raised if the file mode is **IN_FILE**.

Example:

```
CREATE(F, OUT_FILE, "MY_FILE",
      FORM =>
      "WORLD => READ, OWNER => READ_WRITE");
```

The **OPEN** procedure associates an existing external file, of a given **NAME** and **FORM**, with a specified program file object. The procedure also sets the current file mode. If there is an inadmissible change of mode, then the Ada exception **USE_ERROR** is raised.

The **FORM** parameter is a string, formed from a list of attributes, with attributes separated by commas (,). The string is not case sensitive (so that, for example, *HERE* and *here* are treated alike). (FORM attributes are distinct from Ada attributes.) The attributes specify:

- File protection
- File sharing
- File structure
- Buffering
- Appending
- Blocking
- Terminal input

The general form of each attribute is a keyword followed by => and then a qualifier. The arrow and qualifier may sometimes be omitted. The format for an attribute specifier is thus either of

KEYWORD

KEYWORD => QUALIFIER

We will discuss each attribute in turn.

File Protection

These attributes are only meaningful for a call to the CREATE procedure.

File protection involves two independent classifications. The first classification is related to *who* may access the file and is specified by the keywords:

<i>OWNER</i>	Only the owner of the directory may access this file.
<i>GROUP</i>	Only the members of a predefined group of users may access this file.
<i>WORLD</i>	Any user may access this file.

For each type of user who may access a file there are various access *rights*, and this forms the basis for the second classification. In general, there are four types of access right, specified by the qualifiers:

<i>READ</i>	The user may read from the external file.
<i>WRITE</i>	The user may write to the external file.
<i>EXECUTE</i>	The user may execute programs stored in the external file.

NONE

The user has no access rights to the external file. (This access right negates any prior privileges.)

More than one access right may be relevant for a particular file, in which case the qualifiers are linked with underscores (_).

For example, suppose that the *WORLD* may execute a program in an external file, but only the *OWNER* may modify the file.

```
WORLD =>
    EXECUTE ,
OWNER =>
    READ_WRITE_EXECUTE.
```

Repetition of the same qualifier within the attributes is illegal:

```
WORLD =>
    EXECUTE_EXECUTE.                -- NOT legal
```

but repetition of the entire attribute is allowed:

```
WORLD =>
    EXECUTE.
WORLD =>
    EXECUTE.                        -- Legal
```

File Sharing

An external file can be shared, which means associated simultaneously with several logical file objects created by the OPEN and CREATE procedures.

The file sharing attribute may restrict or suppress this capability by specifying one of the following access modes:

NOT_SHARED

Exclusive access - no other logical file may be associated with the external file

SHARED => READERS

Only logical files opened with mode IN are allowed

SHARED => SINGLE_WRITER

Only logical files opened with mode IN and at most one with mode INOUT or OUT are allowed

SHARED => ANY

No restriction

The exception `USE_ERROR` is raised if, for an external file already associated with an Ada file object:

- a further `OPEN` or `CREATE` specifies a file sharing attribute different from the current one
- a further `OPEN`, `CREATE` or `RESET` violates the conditions imposed by the current file sharing attribute.

The restrictions imposed by the file sharing attribute disappear when the last logical file object linked to the external file is closed.

The file sharing attribute provides control over multiple accesses within the program to a given external file.

This control does not extend to the whole system.

The default value for the file sharing attribute is `SHARED => ANY`

File Structure

(a) Text Files

There is no `FORM` parameter to define the structure of text files.

A text file consists of a sequence of bytes holding the ASCII codes of characters.

The representation of Ada-terminators depends on the file's mode (`IN` or `OUT`) and whether it is associated with a terminal device or a mass-storage file:

- Mass-storage files
 - end of line: `ASCII.LF`
 - end of page: `ASCII.LF ASCII.FF`
 - end of file: `ASCII.LF ASCII.EOT`
- Terminal device with mode `IN`
 - end of line: `ASCII.LF`
 - end of page: `ASCII.LF ASCII.FF`
 - end of file: `ASCII.LF ASCII.FF`

- Terminal device with mode OUT

end of line: ASCII.LF
 end of page: ASCII.FF
 end of file: ASCII.EOT

(b) Binary Files

Two FORM attributes, *RECORD_SIZE* and *RECORD_UNIT*, control the structure of binary files.

A binary file can be viewed as a sequence (sequential access) or a set (direct access) of consecutive RECORDS.

The structure of such a record is:

[HEADER] OBJECT [UNUSED_PART]

and it is formed from up to three items:

- an OBJECT with the exact binary representation of the Ada object in the executable program, possibly including an object descriptor
- a HEADER consisting of two fields (each of 32 bits):
 - the length of the object in bytes
 - the length of the descriptor in bytes
- an UNUSED_PART of variable size to permit full control of the record's size

The HEADER is implemented only if the actual parameter of the instantiation of the IO package is unconstrained.

The file structure attributes take the form:

RECORD_SIZE => *size_in_bytes*

RECORD_UNIT => *size_in_bytes*

Their meaning depends on the object's type (constrained or not) and the file access mode (sequential or direct access):

- If the object's type is constrained:
 - The *RECORD_UNIT* attribute is illegal
 - If the *RECORD_SIZE* attribute is omitted, no UNUSED_PART will be implemented: the default *RECORD_SIZE* is the object's size

- If present, the *RECORD_SIZE* attribute must specify a record size greater than or equal to the object's size, otherwise the exception *USE_ERROR* will be raised
- b) If the object's type is unconstrained and the file access mode is direct:
- The *RECORD_UNIT* attribute is illegal
 - The *RECORD_SIZE* attribute has no default value, and if it is not specified, a *USE_ERROR* will be raised
 - An attempt to input or output an object larger than the given *RECORD_SIZE* will raise the exception *DATA_ERROR*
- c) If the object's type is unconstrained and the file access mode is sequential:
- The *RECORD_SIZE* attribute is illegal
 - The default value of the *RECORD_UNIT* attribute is 1 (byte)
 - The record size will be the smallest multiple of the specified (or default) *RECORD_UNIT* that holds the object and its length. This is the only case where records of a file may have different sizes.

Buffering

The buffer size can be specified by the attribute

BUFFER_SIZE => *size_in_bytes*

The default value for *BUFFER_SIZE* is 0 (which means no buffering) for terminal devices; it is 1 sector for disk files.

Appending

Only to be used with the procedure *OPEN*, the format of this attribute is simply

APPEND

and it means that any output will be placed at the end of the named external file.

In normal circumstances, when an external file is opened, an index is set which points to the beginning of the file. If the *APPEND* attribute is present for a sequential or for a text file, then data transfer will commence at the end of the file. For a direct access file, the value of the index is set to one more than the number of records in the external file.

This attribute is not applicable to terminal devices.

Blocking

This attribute has two alternative forms:

BLOCKING .

or

NON_BLOCKING .

This attribute specifies the IO system behavior desired at any moment that a request for data transfer cannot be fulfilled. The stoppage may be due, for example, to the unavailability of data, or to the unavailability of the external file device.

NON_BLOCKING

If this attribute is set, then the task that ordered the data transfer is suspended - meaning that other tasks can execute. The suspended task is kept in a 'ready' state, together with other tasks in a ready state at the same priority level (that is, it is rescheduled).

When the suspended task is next scheduled, the data transfer request is reactivated. If ready, the transfer is activated, otherwise the rescheduling is repeated. Control returns to the user program after completion of the data transfer.

BLOCKING

In this case the task waits until the data transfer is complete, and all other tasks are suspended (or 'blocked'). The system is busy waiting.

The default for this attribute depends on the actual program: it is *BLOCKING* for programs without task declarations and *NON_BLOCKING* for a program containing tasks.

Terminal input

This attribute takes one of two alternative forms:

TERMINAL_INPUT => *LINES*.

TERMINAL_INPUT => *CHARACTERS*.

Terminal input is normally processed in units of a line at a time, where a line is delimited by a special character. A process attempting to read from the terminal as an external file will be suspended until a complete line has been typed. At that time, the outstanding read call (and possibly also later calls) will be satisfied.

The first option specifies line-at-a-time data transfer, which is the default case.

The second option means that data transfer is character by character, and so a complete line does not have to be entered before the read request can be satisfied. For this option the `BUFFER_SIZE` must be zero.

The `TERMINAL_INPUT` attribute is only applicable to terminal devices.

9 Characteristics of Numeric Types

9.1 Integer Types

The ranges of values for integer types declared in package `STANDARD` are as follows:

<code>SHORT_SHORT_INTEGER</code>	-128 .. 127	-- $2^{**7} - 1$
<code>SHORT_INTEGER</code>	-32768 .. 32767	-- $2^{**15} - 1$
<code>INTEGER</code>	-2147483648 .. 2147483647	-- $2^{**31} - 1$

For the packages `DIRECT_IO` and `TEXT_IO`, the range of values for types `COUNT` and `POSITIVE_COUNT` are as follows:

<code>COUNT</code>	0 .. 2147483647	-- $2^{**31} - 1$
<code>POSITIVE_COUNT</code>	1 .. 2147483647	-- $2^{**31} - 1$

For the package `TEXT_IO`, the range of values for the type `FIELD` is as follows:

<code>FIELD</code>	0 .. 255	-- $2^{**8} - 1$
--------------------	----------	------------------

9.2 Floating Point Type Attributes

	<code>FLOAT</code>	<code>LONG_FLOAT</code>
<code>DIGITS</code>	6	15
<code>MANTISSA</code>	21	51
<code>EMAX</code>	84	204
<code>EPSILON</code>	9.53674E-07	8.88178E-16
<code>LARGE</code>	1.93428E+25	2.57110E+61
<code>SAFE_EMAX</code>	125	1021
<code>SAFE_SMALL</code>	1.17549E-38	2.22507E-308

SAFE_LARGE	4.25353E+37	2.24712E+307
FIRST	-3.40282E+38	-1.79769E+308
LAST	3.40282E+38	1.79769E+308
MACHINE_RADIX	2	2
MACHINE_EMAX	128	1024
MACHINE_EMIN	-125	-1021
MACHINE_ROUNDS	true	true
MACHINE_OVERFLOWS	false	false
SIZE	32	64

9.3 Attributes of Type DURATION

DURATION'DELTA	0.001
DURATION'SMALL	0.0009765625 (= 2**(-10))
DURATION'FIRST	-2097152.0
DURATION'LAST	2097151.999
DURATION'LARGE	same as DURATION'LAST

10 Other Implementation-Dependent Characteristics

10.1 Use of the Floating-Point Coprocessor (80287, 80387)

The Alsys 386 Unix Ada Compiler generates instructions to use the floating point coprocessor for all floating point operations (but, of course, not for operations involving only *universal_real*).

A floating point coprocessor, 80287 or 80387, is required for the execution of programs that use arithmetic on floating point values. The coprocessor is needed if the FLOAT_IO or FIXED_IO packages of TEXT_IO are used.

The *Runtime Executive* will detect the absence of the floating point coprocessor if it is required by a program and will raise NUMERIC_ERROR.

10.2 Characteristics of the Heap

UNCHECKED_DEALLOCATION is implemented for all Ada access objects except access objects to tasks. Use of UNCHECKED_DEALLOCATION on a task object will lead to unpredictable results.

All objects whose visibility is linked to a subprogram, task body, or block have their storage reclaimed at exit.

The maximum size of the heap is limited only by available memory. This includes the amount of physical memory (RAM) and the amount of virtual memory (hard disk swap space).

All objects created by allocators go into the heap. Also, portions of the *Runtime Executive* representation of task objects, including the task stacks, are allocated in the heap.

10.3 Characteristics of Tasks

The default task stack size is 1K bytes (32K bytes for the environment task), but by using the Binder option STACK.TASK the size for all task stacks in a program may be set to a size from 1K bytes to 64K bytes.

Normal priority rules are followed for preemption, where PRIORITY values are in the range 1 .. 10. A task with *undefined* priority (no pragma PRIORITY) is considered to be lower than priority 1.

The minimum timeable delay is 1.0/1000. seconds. This is the finest resolution provided by 386 UNIX.

The maximum number of active tasks is restricted only by memory usage.

The acceptor of a rendezvous executes the accept body code in its own stack. Rendezvous with an empty accept body (for synchronization) does not cause a context switch.

The main program waits for completion of all tasks dependent upon library packages before terminating.

Abnormal completion of an aborted task takes place immediately, except when the abnormal task is the caller of an entry that is engaged in a rendezvous, or if it is in the process of activating some tasks. Any such task becomes abnormally completed as soon as the state in question is exited.

The message

GLOBAL BLOCKING SITUATION DETECTED

is printed to STANDARD_OUTPUT when the *Runtime Executive* detects that no further progress is possible for any task in the program. The execution of the program is then abandoned.

10.4 Definition of a Main Subprogram

A library unit can be used as a main subprogram if and only if it is a procedure that is not generic and that has no formal parameters.

10.5 Ordering of Compilation Units

The Alsys 386 Unix Ada Compiler imposes no additional ordering constraints on compilations beyond those required by the language.

11 Limitations

11.1 Compiler Limitations

- The maximum identifier length is 255 characters.
- The maximum line length is 255 characters.
- The maximum number of unique identifiers per compilation unit is 2500.

11.2 Hardware Related Limitations

- The maximum size of the generated code for a single compilation unit is 65535 bytes.
- The maximum size of a single array or record object is 65522 bytes. The maximum size of a static record is 4096 bytes.
- The maximum size of a single stack frame is 32766 bytes, including the data for inner package subunits unnested to the parent frame.
- The maximum amount of data in the global data area is 65535 bytes, including compiler generated data that goes into the GDA (about 8 bytes per compilation unit plus 4 bytes per externally visible subprogram).
- The maximum amount of data in the heap is limited only by available memory, real and virtual.

INDEX

80287 38

80387 38

Abnormal completion 39

Aborted task 39

Allocators 39

Application Developer's Guide 2

Array objects 28

Array subtype 4

Array type 4

ASSIGN_TO_ADDRESS 7

Attributes of type DURATION 38

Binder 39

Characteristics of tasks 39

Compiler limitations 40

 maximum identifier length 40

 maximum line length 40

 maximum number of compilation
 units 40

 maximum number of unique
 identifiers 40

COUNT 37

DIGITS 37

DIRECT_IO 37

DURATION'DELTA 38

DURATION'FIRST 38

DURATION'LARGE 38

DURATION'LAST 38

DURATION'SMALL 38

EMAX 37

Empty accept body 39

EPSILON 37

FETCH_FROM_ADDRESS 7

FIELD 37

FIRST 38

FIXED_IO 38

FLOAT_IO 38

Floating point coprocessor 38

Floating point operations 38

Floating point type attributes 37

GLOBAL BLOCKING SITUATION
DETECTED 39

Hardware limitations

 maximum amount of data in the
 global data area 40

 maximum data in the heap 40

 maximum size of a single array or
 record object 40

 maximum size of a single stack frame
 40

 maximum size of the generated code
 40

Hardware related limitations 40

Heap 39

Implementation generated names 28

INTEGER 37

Integer types 37

Intel object module format 3

INTERFACE 2, 3

INTERFACE_NAME 2, 3

LARGE 37

LAST 38

Library unit 40

Limitations 40

MACHINE_EMAX 38

MACHINE_EMIN 38

MACHINE_MANTISSA 38

MACHINE_OVERFLOWS 38

MACHINE_RADIX 38

MACHINE_ROUNDS 38

Main program 39

Main subprogram 40

MANTISSA 37

Maximum amount of data in the global
 data area 40

Maximum data in the heap 40

Maximum identifier length 40

Maximum line length 40

Maximum number of compilation units
 40

Maximum number of unique identifiers
 40

- Maximum size of a single array or record object 40
- Maximum size of a single stack frame 40
- Maximum size of the generated code 40
- Microsoft Linkers 3
- Minimum timeable delay 39
- Number of active tasks 39
- NUMERIC_ERROR 38
- Ordering of compilation units 40
- P'IS_ARRAY 4
- Parameter passing 1
- POSITIVE_COUNT 37
- Pragma IMPROVE 4
- Pragma INDENT 3
- Pragma INTERFACE 2, 3
- Pragma INTERFACE_NAME 2, 3
- Pragma PACK 4
- Pragma PRIORITY 4, 39
- Pragma SUPPRESS 4
- Predefined packages 28
- PRIORITY 4, 39
- Record objects 28
- Rendezvous 39
- Runtime Executive 1, 3, 38, 39
- SAFE_EMAX 37
- SAFE_LARGE 38
- SAFE_SMALL 37
- SHORT_INTEGER 37
- SHORT_SHORT_INTEGER 37
- SIZE 38
- STANDARD_OUTPUT 39
- Storage reclamation at exit 39
- SUPPRESS 4
- SYSTEM 4
- Task stack size 39
- Task stacks 39
- Tasks
 - characteristics of 39
- TEXT_IO 37
- Unchecked conversions 29
- UNCHECKED_DEALLOCATION 39
- Universal_real 38
- Unix Linker 3